

MOHAMMAD ALI JINNAH UNIVERSITY

**CONVERGENCE OF INTERNET PROTOCOL
WITH
THIRD GENERATION MOBILE NETWORKS**

A path to fourth generation networks

Submitted By:

Ahmed Wajih and Faisal Haider (FA02-BT-0052)

Supervised By:

**Professor Mr. Tabish Rizvi
Department of Computer Science**

**IMPLEMENTATION OF INTERNET PROTOCOL
FOR
THIRD GENERATION NETWORKS**

A path to fourth generation networks

Submitted By:

Ahmed Wajih and Faisal Haider (FA02-BT-0052)

A Research Project in partial fulfillment of the requirements for the degree of
Bachelor of Science in Computer Science
BS (CS)

MOHAMAD ALI JINNAH UNIVERSITY

Acknowledgments

First and foremost, we would like to thank our Creator, the most Beneficent the most Merciful, for blessing us with something that many are deprived from – Education. We have been given a great chance in our life to learn and gain knowledge and then apply it elsewhere.

Secondly, we want to acknowledge the role of our parents in our life, our education in particular. We simply couldn't have achieved all this without their help, support and understanding each and every stage of our learning and life.

Next, we want to thank our teachers, especially Mr. Waqar Rizvi and Mr. Shua Hussain who taught and guided us really well to achieve our goal – and Mr. Tabish Rizvi, the advisor of our project. We thank them for all their time.

Lastly, we want to thank our family, The most.

Chapter 1 Introduction

1.2 IP	1
1.2.1 3G	2
1.2.2 IP for 3G	2
1.3 Engineering Reasons for 'IP for 3G'	3
1.3.1 IP Design Principles	3
1.3.2 Benefits of the IP approach	5
1.3.3 Weaknesses of the IP approach	5
1.4 Economic Reasons for 'IP for 3G'	6
1.4.1 3G Business Case	7
1.4.2 Impact of 'IP for 3G' on Business Case	12
1.5 Conclusion	13

Chapter 2 Introduction to 3G..... **14**

2.1 Introduction to 3G	14
2.2 Mobile Standards	15
2.2.1 Who's who in 3G Standards	16
2.3 History of 3G	17
2.3.1 Pre-1996 - The Research Trimester	17
2.3.2 1996–1998 - The IMT 2000 Trimester	19
2.3.3 1998 Onwards - The Standardization Trimester	21
2.4 Spectrum - The 'Fuel' of Mobile Systems	22

Chapter 3 UMTS **25**

Network Overview	25
3.1 UMTS Architecture ..	28
3.2 UMTS Security	30
3.3 UMTS Communication Management	32
3.3.1 Connection Management	32
3.3.2 Session Management	32
3.4 UMTS QoS	34
3.5 UMTS Mobility Management	37
3.6 UMTS Core Network Transport	37
3.7 Signaling in the UMTS Core Network	40

Chapter 4 **UTRAN** 41

<i>UMTS Radio Access Network (UTRAN)</i>	41
<i>4.1 The W-CDMA Air Interface and the Uu Interface</i>	41
<i>4.2 UTRAN Mobility Management</i>	43
4.2.1 <i>Soft Handover</i>	43
4.2.2 <i>Handover Types</i>	44
4.2.3 <i>UTRAN Transport</i>	45
4.2.4 <i>UTRAN QoS</i>	47
4.2.5 <i>UTRAN Signaling</i>	49
<i>4.3 cdma2000 Packet Core Network</i>	49
<i>4.4 Conclusion</i>	52

Chapter 5 **Session Management** 54

<i>5.1 Introduction</i>	54
<i>5.2 Session Management</i>	54
5.2.1 <i>What is a Session?</i>	54
5.2.2 <i>Functions of Session Management Protocols</i>	55
5.2.3 <i>Summary</i>	56
<i>5.3 Current Status</i>	56
5.3.1 <i>Session Management</i>	56
5.3.2 <i>Within 2G Networks</i>	57
5.3.3 <i>Within the Internet</i>	57
5.3.4 <i>Session Management for Future Applications</i>	58
<i>5.4 VHE Concept</i>	58
5.4.1 <i>Within 2G/3G Networks</i>	58
5.4.2 <i>Future VHE</i>	59
<i>5.5 Session Initiation Protocols</i>	60
5.5.1 <i>H.323</i>	60
5.5.2 <i>SIP</i>	60
5.5.3 <i>Session Initiation for 3G</i>	61
<i>5.6 SIP in Detail</i>	61
5.6.1 <i>Basic Operation of SIP</i>	61
5.6.2 <i>SIP and User Location</i>	63
5.6.3 <i>Characteristics of SIP</i>	64
<i>5.7 SIP in Use</i>	65
5.7.1 <i>Connecting IP and Telephony</i>	65
5.7.2 <i>SIP Supported Services</i>	66
<i>5.8 Conclusions</i>	68
5.8.1 <i>SIP</i>	68
5.8.2 <i>VHE</i>	69

Chapter 6	IP Mobility	71
6.1 Scope		71
6.2 Introduction - What is IP Mobility?		71
6.2.2 The Problem of IP Mobility		72
6.2.3 Layer 2 Solutions		72
6.2.4 Application-layer Solution		73
6.2.5 Layer 3 Solutions		73
6.2.6 Locators vs. Identifiers		73
6.3 SIP - A Protocol for Personal Mobility		75
6.4 Introduction to Terminal Mobility		76
6.4.1 Macromobility vs. Micromobility		77
6.5 Mobile IP - A Solution for Terminal Macromobility		78
6.5.1 Outline of Mobile IP		78
6.5.2 Mobile IPv4		80
6.5.3 Triangular Routing and Route Optimization		80
6.5.4 Reverse Tunneling		80
6.5.5 NAT Traversal		81
6.5.6 Address shortage		81
6.5.7 Foreign Agents		81
6.5.8 Mobile IPv6		81
6.5.9 Relationship of SIP and Mobile IP		82
6.6 Terminal Micromobility		83
6.6.1 Introduction		83
6.7 Mobile IP-based schemes		84
6.7.1 Local Mobility Agents		84
6.7.2 'Fast and Smooth' Mobile IP-based Schemes		84
6.7.3 Per-host Forwarding Schemes		84
6.7.4 New Schemes		84
6.7.5 MANET-based Schemes		84
6.7.6 Multicast-based Schemes		85
6.8 Mobile IP-based Protocols		85
6.8.1 Local Mobility Agent Schemes		85
6.8.2 Regional Registration for Mobile IPv4		86
6.8.3 Home Registration		86
6.8.4 Regional Registration Request		87
6.8.5 Hierarchical Mobile IPv6		87
6.8.6 'Fast and Smooth' Mobile IP-based Schemes		88
6.8.7 Two CoAs		88
6.8.8 Simultaneous Bindings and Packet Bi-casting		89
6.8.9 Temporary Tunnel		89
6.8.10 Fast Handovers for Mobile IPv6		89
6.8.11 Fast Handovers for Mobile IPv4		90
6.9 Per-host Forwarding Protocols		91
6.9.1 Outline of their Operation		91
6.9.2 New Protocols - Cellular IP and Hawaii		92

6.9.3 Cellular IPv4	92
6.9.4 Cellular IPv6	95
6.9.5 Multicast-based Schemes	98

Chapter 7 Comparison of Micromobility Protocols

<i>Comparison of Micromobility Protocols</i>	100
7.1 Comparison of Micromobility Protocols	100
7.2 Operation	100
7.3 Architecture	101
7.3.1 Handover can be Initiated Using Layer 2 'Triggers'	101
7.3.2 Handover is Mobile-controlled	101
7.3.3 Packet Bi-casting and Buffering	102
7.3.4 Separate 'Handover' from 'Path Updates'	102
7.3.5 Separate Paging Protocol	103
7.3.6 Separate Macro and Micro Mobility	103
7.4 Scalability	104
7.4.1 Tunnel-based Schemes	104
7.4.2 Per-host Mobility Schemes	104
7.4.3 Soft vs. Hard State	106
7.4.4 Address Allocation	106
7.4.5 Multiple Access Network Gateways	106
7.4.6 Paging	106
7.5 Reliability	107
7.5.1 Tunnel-based Mobility Schemes	107
7.5.2 Per-host Schemes	108
7.6 Philosophy	108
7.6.1 QoS Compatibility	109
7.6.2 Web Caching	110
7.7 Other Aspects of Terminal Mobility	110
7.8. Context (or State) Transfer	111
7.9 Paging and Dormant Mode Management	112
7.9.1 Paging	114
7.9.2 Within the Access Network	114
7.9.3 Over the Air	114
7.10 A Brief Word on Security for Mobility Management	115
7.11 Conclusions	116

Chapter 8 Quality of Service 119

<i>Quality of Service</i>	119
8.1 Introduction	119
8.1.1 What is QoS?	119
8.1.2 Cost/Complexity/Strength of QoS Compromise	120
8.1.3 QoS Co-operation	121
8.2 Current IP QoS Mechanisms	121
8.2.2 Random Early Detect and Explicit Congestion Notification	125
8.2.3 Basic RTP	125
8.2.4 Conclusions	128
8.3. Functionality Required of the Network to Support QoS	128
8.3.1 Loss Management	130
8.3.2 Low signal-to-noise ratio	130
8.3.3 Mechanisms to Provide Network QoS	132
8.3.4 Prioritization and Reservation	133
Characteristics of Signalling Systems	133
Wireless Efficiency	134
8.3.5 Admission Control	135
8.3.6 Call Admission Control Architectures	135
8.3.7 Admission Control Descriptions	136
8.3.8 Traffic Classification and Conditioning	136
8.3.9 Mobility Issues	137
8.3.10 QoS Management during Handover	137
8.3.11 Context Transfer Protocol	139
8.3.12 QoS Management after Handover	139
8.4.1 IntServ	140
8.4.2 DiffServ	142
8.4.3 RSVP	144
8.4.4 Details of RSVP Signalling	144
8.4.5 Use of RSVP in a Mobile Environment	145
8.5 IP QoS for 3G - A Possible Solution	147
8.5.2 Bounded Delay Differentiated Service	150
8.5.3 Basic Operation of Bounded Delay Service	150
8.5.4 Signalling	152
8.6 Conclusions	155

Chapter 9	IP for 3G	156
9.1 Introduction		156
9.2.1 Principles		157
9.2.2 Routing and Mobility		159
9.2.3. Security		161
9.2.4 An Answer		163
9.3 Advantages of an All-IP Network		164
9.4 3G Network Evolution		166
9.4.2 UMTS R5 - IP Call Control and Signalling		167
Call Control		168
Roaming		169
IPv6		171
9.4.4 CDMA2000 Evolution		172
9.5 UMTS Beyond R5		173
9.6 Wireless LANs		174

Chapter 1

Introduction

1.1 IP for 3G

This section concerns 'IP for 3G' and explains what is meant by the terms 'IP' and '3G'. It also hopefully positions it with regard to things that readers may already know about IP or 3G, i.e. previous knowledge is helpful but not a prerequisite.

1.2 IP

What is meant by 'IP' in the context of this Report?

IP stands for the 'Internet Protocol', which specifies how to segment data into packets, with a header that (amongst other things) specifies the two end points between which the packet is to be transferred. 'IP' in the context of this report should not be interpreted in such a narrow sense, but rather more generally as a synonym for the 'Internet'.

The word 'Internet' has several connotations. First, and most obviously, 'Internet' refers to 'surfing' the user's activity of looking at web pages, ordering goods on-line, doing e-mail and so on, which can involve accessing public sites or private (internal company) sites. This whole field of applications and the user experience are not the focus of this report. Instead, attention is focused on the underlying network and protocols that enable this user experience and such a range of applications. Next, 'Internet' refers to the network, i.e. the routers and links over which the IP packets generated by the application (the 'surfing') are transferred from the source to the destination.

Then, there are the 'Internet' protocols - the family of protocols that the Internet network and terminal run; things like TCP (Transmission Control Protocol, which regulates the source's transmissions) and DHCP (Dynamic Host Configuration Protocol, which enables terminals to obtain an IP address dynamically).

The term 'Internet' can also be used more loosely to refer to the IETF - the Internet Engineering Task Force, which is the body that standardizes Internet protocols. It is noteworthy for its standardization process being: (1) open - anyone can contribute (for free) and attend meetings; (2) pragmatic - decisions are based on rough consensus and running code

The Internet standardization process appears to be faster and more dynamic than that of traditional mobile standardization organizations - such as ETSI, for example. However, in reality, they are trying to do rather different jobs. In the IETF, the emphasis is on protocols one protocol per function (thus, TCP for transport, HTTP for hypertext transport and so forth). The IETF has only a very loose architecture and general architectural principles. Many details of building IP systems are left to integrators and manufacturers. In contrast, the standards for GSM, for example, are based around a fixed architecture and tightly defined interfaces (which include protocols). The advantage of defining interfaces, as opposed to just protocols, is that that much more of the design work has been done and equipment from different manufactures will always inter-

operate. As will be seen later, there is a large amount of work to be done to turn the IETF protocols into something that resembles a mobile architecture, and introduces some fixed elements and interfaces to accomplish this.

Finally, 'Internet' can also imply the 'design principles' that are inherent in the Internet protocols.

1.2.1 3G

What is meant by '3G' in the context of this Report?

'3G' is short for 'third generation mobile systems'. 3G is the successor of 2G - the existing digital mobile systems: GSM in most of the world, D-AMPS in the US, and PHS and PDC in Japan. 2G in turn was the successor of 1G -the original analogue mobile systems. Just as for 'IP', the term '3G' also has several connotations.

First, '3G' as in its spectrum: the particular radio frequencies in which a 3G system can be operated. 3G has entered the consciousness of the general public because of the recent selling off of 3G spectrum in many countries and, in particular, the breathtaking prices reached in the UK and Germany. From a user's perspective, '3G' is about the particular services it promises to deliver. 1G and 2G were primarily designed to carry voice calls; although 2G's design also includes 'short message services', the success of text messaging has been quite unexpected. 3G should deliver higher data rates (up to 2 Mbit/s is often claimed, though it is likely to be much lower for many years and in many environments), with particular emphasis on multimedia (like video calls) and data delivery.

The term '3G' also covers two technical aspects. First is the air interface, i.e. the particular way in which the radio transmission is modulated in order to transfer information 'over the air' to the receiver. For most of the 3G systems being launched over the next few years, the air interface is a variant of W-CDMA (Wideband Code Division Multiple Access). The second technical aspect of '3G' is its network. The network includes all the base stations, switches, gateways, databases and the (wired) links between them, as well as the definition of the interfaces between these various components (i.e. the architecture). Included here is how the network performs functions such as security (e.g. authenticating the user), quality of service (e.g. prioritizing a video call over a data transfer) and mobility management (e.g. delivering service when moving to the coverage of an adjacent base station). Several specific 3G systems have been developed, including UMTS in Europe and cdma2000 in the US. A reasonable summary is that the 3G network is based on an evolved 2G network.

1.2.2 IP for 3G

What is meant by IP for 3G? 3G systems will include IP multimedia allowing the user to browse the Internet, send e-mails, and so forth. There is also a second phase of UMTS being developed, as will be discussed later, that specifically includes something called the Internet Multimedia Subsystem. Why, then, is IP argued for in 3G? The issue of IP for 3G is really more about driving changes to Internet protocols to make them suitable to provide 3G functionality supporting aspects like handover of real-time services and guaranteed QoS. If a 3G network could be built using (enhanced) IP routers and servers and common IP protocols, then

- It might be cheaper to procure through economies of scale due to a greater commonality with fixed networks.
- It could support new IP network layer functionality, such as multicast and any-cast, natively, i.e. more cheaply without using bridges, etc.
- It would offer operators greater commonality with fixed IP networks and thus savings from having fewer types of equipment to maintain and the ability to offer common fixed/mobile services.
- It would be easier for operators to integrate other access technologies (such as wireless LANs) with wide-area cellular technologies.

So, IP for 3G is about costs and services; if IP mobility, QoS, security and session negotiation protocols can be enhanced/developed to support mobile users, including 3G functionality such as real-time handover, and a suitable IP architecture developed, then we believe there will be real benefits to users and operators. This Report, then, is largely about IP protocols and how current research is moving in these areas. The final chapter attempts to build an architecture that uses native IP routing and looks at how some of this functionality is already being included in 3G standards.

1.3 Engineering Reasons for 'IP for 3G'

Here, only preliminary points are outlined basically providing some hints as to why the Report covers the topics it does and where it is going. One way into this is to examine the strengths and weaknesses of IP and 3G. The belief, therefore, is that 'IP for 3G' would combine their strengths and alleviate their weaknesses. At least it indicates the areas that research and development need to concentrate on in order for 'IP for 3G' to happen.

1.3.1 IP Design Principles

Perhaps the most important distinction between the Internet and 3G (or more generally the traditional approach to telecoms) is to do with how they go about designing a system. There are clearly many aspects involved; security, QoS, mobility management, the service itself, the link layer technology (e.g. the air interface), the terminals, and so on. The traditional telecoms approach is to design everything as part of a single process, leading to what is conceptually a single standard (in reality, a tightly coupled set of standards). Building a new system will thus involve the design of everything from top to bottom from scratch (and thus it is often called the 'Stovepipe Approach'). By contrast, the IP approach is to design a 'small' protocol that does one particular task, and to combine it with other protocols (which may already exist) in order to build a system. IP therefore federates together protocols selected from a loose collection. To put it another way, the IP approach is that a particular layer of the protocol stack does a particular task. This is captured by the IP design principle, *always keep layer transparency*, or by the phrase, *IP over everything and everything over IP*. This means that IP can run on top of any link layer (i.e. bit transport) technology and that any service can run on top of IP. Most importantly, the service is not concerned with, and has no knowledge of, the link layer. The analogy is often drawn with the hourglass, e.g., with its narrow waist representing the simple, single IP layer ([Figure 1.1](#)). The key requirement is to have a well-defined interface between the layers, so that the layer above knows what behavior to expect from the layer below, and what functionality it can use. By contrast, the Stovepipe

Approach builds a vertically integrated solution, i.e. the whole system, from services through network to the air interface, is designed as a single entity. So, for example in 3G, the voice application is specially designed to fit with the W-CDMA air interface.

INTRODUCTION




Figure 1.1: IP over everything and everything over IP. The Internet's 'hourglass' protocol stack..

Another distinction between the Internet and 3G is where the functionality is placed. 3G (and traditional telecoms networks) places a large amount of functionality within the network, for example at the Mobile Switching Centre. The Internet tries to avoid this, and to confine functionality as far as possible to the edge of the network, thus keeping the network as simple as possible. This is captured by the IP design principle: *always think end to end*.

It is an assertion that the end systems (terminals) are best placed to understand what the applications or user wants. The principle justifies why IP is connectionless (whereas the fixed and mobile telephony networks are connection-oriented). So, every IP packet includes its destination in its header, whereas a connection-oriented network must establish a connection in advance, i.e. before any data can be transferred. One implication is that, in a connection oriented network, the switches en route must remember details of

the connection (it goes between this input and that output port, with so much bandwidth, and a particular service type, etc.).

1.3.2 Benefits of the IP approach

IP is basically a connectionless packet delivery service that can run over just about any Layer 2 technology. In itself, it is not the World Wide Web or e-mail or Internet banking or any other application. IP has been successful because it has shown that for non-real-time applications, a connectionless packet service is the right network technology. It has been helped by the introduction of optical fiber networks, with their very low error rates, making much of the heavyweight error correction abilities of older packet protocols like X25 unnecessary.

IP also decouples the network layer very clearly from the service and application. Operating systems like Windows have IP sockets that can be used by applications written by anyone; a lone programmer can devise a new astrology calculator and set up a server in his garage to launch the service. Because IP networks provide so little functionality (IP packet delivery), the interfaces to them are simple and can be opened without fear of new services bringing the network down, the point being that IP connectivity has become a commodity and it has been decoupled (by the nature of IP) from the content/applications.

IP applications also tend to make use of end-to-end functionality: when a user is online to their bank, they require that their financial details be heavily encrypted. This functionality could have been provided by the network, but instead, it is done on a secure sockets layer above the IP layer in the browser and the bank's server. Clearly, this is a more flexible approach the user can download a certificate and upgrade to 128-bit security instantly if the network were providing the service, there would be a requirement for signaling, and new features would have to be integrated and tested with the rest of the features of the network.

1.3.3 Weaknesses of the IP approach

IP is not a complete architecture or a network design it is a set of protocols. If a number of routers were purchased and connected to customers, customers could indeed be offered a connectionless packet delivery service. It would quickly become apparent that the amount of user traffic entering your network would need to be limited (perhaps through charging). To make sure that everybody had a reasonable throughput, the network would have to be over provisioned. A billing engine, network management platform (to identify when the routers and connections break), and help desk would be needed also, in other words, quite a lot of the stuff of a more 'traditional' fixed network.

If customers then said that they wanted real-time service support (to run voice, say), something like an ATM network underneath the IP would need to be installed, to guarantee that packets arrive within a certain maximum delay. In fact, IP is fundamentally unsuited to delivering packets within a time limit and, as will be seen in [later](#), adding this functionality, especially for mobile users, is a very hot IP research topic. In the end, adding real-time QoS to IP will mean 'fattening' the hourglass and losing some of the simplicity of IP networks.

IP networks also rely on the principle of global addressing, and this IP address is attached to every packet. Unfortunately, there are not enough IP addresses to go round; since the address field is limited to 32 bits. Consequently, a new version of the IP protocol - IPv6 is being introduced to extend the address space to 128 bits. The two versions of IP also have to sit in the hourglass fattening it still further.

Another issue is that the Internet assumes that the end points are fixed. If a terminal moves to a new point of attachment, it is basically treated in the same as a new terminal. Clearly, a mobile voice user, for example, will expect continuous service even if they happen to have handed over, i.e. moved on to a new base station. Adding such mobility management functionality is another key area under very active investigation.

Because IP connectivity is just a socket on a computer, it is quite often the case that applications on different terminals are incompatible in some way - there is no standard browser, as some people use Netscape, some use Internet Explorer, some have version 6, and so forth. When browsing, this is not too much trouble, and the user can often download new plug-ins to enhance functionality. When trying to set up something like a real-time voice call, however, this means quite a lot of negotiation on coding rates and formats, etc. In addition, the user's IP address will change at each log in (or periodically on DSL supported sessions also) - meaning that individuals (as opposed to servers using DNS) are nearly impossible to locate instantly for setting up a voice session. What is needed in IP is a way of identifying users that is fixed (e.g. comparable with an e-mail address), binding it more rapidly to one (or more) changing IP addresses, and then being able to negotiate sessions (agreeing such things as coding rates and formats). Later we will provide details on how the Session Initiation Protocol (SIP) is able to fulfill this role.

It is interesting that some of the approaches to solving these downsides involve 'weakening' our two IP design principles; for example by adding quality-of-service state to some routers (i.e. weakening the end-to-end principle) or adding inter-layer hints between the link and IP layers (e.g. radio power measurements are used to inform the IP layer that a handover is imminent, i.e. weakening the layer transparency principle). So, a key unanswered question is: to what extent should the IP design principles which have served the Internet so well be adapted to cope with the special problems of wireless-ness and mobility?.

1.4 Economic Reasons for 'IP for 3G'

As already indicated, IP for 3G is about reducing costs. There is nothing that IP for 3G will enable that cannot already be done in 3G at a price. IP is just a connectionless packet delivery service, and a 3G network could be thought of as a Layer 2 network. The Layer 2 (3G) might not support multicast, but that can still be emulated with a series of point-to-point connections. What adoption of IP protocols and design principles might do for 3G is reduce costs; this section delves deeper into exactly where 3G costs arise and explains in detail how an IP-based evolution could, potentially, reduce them.

1.4.1 3G Business Case

3G Costs

First, there is the cost of the spectrum. This varies wildly from country to country (see [Table 1.1](#)) from zero cost in Finland and Japan, up to \$594 per capita in Britain.

Table 1.1: Licence cost (\$) per capita in selected countries	
Country	Cost per capita (US\$)
UK	594.20
Germany	566.90
Italy	174.20
Taiwan	108.20
US	80.90
South Korea	60.80
Singapore	42.60
Australia	30.30
Norway	20.50
Switzerland	16.50
Spain	11.20
Sweden	5.70
Japan	0.00
Finland	0.00
Note: US auction was for PCS Licences that can be upgraded later to 3G.	
Source: 3G Newsroom [3] .	

Second, there is the cost of the 3G network itself the base stations, switches, links, and so on. It is higher than for a 2G network, because the base station sites need to be situated more densely, owing to the frequency of operation and the limited spectrum being used to support broadband services. For example, the consultancy Ovum estimates the cost as more than \$100 billion over the next five years in Europe alone, whereas for the UK, Crown Castle estimate that a 3G operator will spend about £2850 million on infrastructure (i.e. capital expenditure) with an annual operating cost of £450 million (including: £840 million on sites; £1130 million on Node Bs, £360 million on RNCs; £420 million on backhaul and £100 million on the Core Network)

These large amounts are a strong incentive for 3G operators to try to find ways of sharing infrastructure and so share costs. For example, Mobil-COM (a German operator) estimates that 20-40% can be saved, mainly through collocating base stations ('site sharing'), and in our UK example, Crown Castle argues that the capital spend can be cut by almost one-third to £2 billion . However, sharing may not be in the interests of all operators Ovum outlines some of the pros and cons depending on the operator's market position but the burst of the dot.com bubble and the global economic downturn have certainly increased interest in the idea. Infrastructure sharing may not be permitted in all

countries for example, the conditions attached to a license may not allow it but regulators are being increasingly flexible (e.g. UK, France). Some governments (e.g. the French and Spanish) are also reducing the license cost from the agreed amount.

3G Services and Income

A large number of services have been suggested for 3G. Here, we look at a few of them.

Lessons from 2G - Voice

2G systems like GSM and D-AMPS have shown that voice communication is a very desirable service and that customers will pay a considerable premium for the advantage of mobility a combination of being reachable anywhere anytime and having one's own personal, and personalized, terminal. For any 3G operator who does not have a 2G license, voice will of course be a very important service. But for all operators, it is likely to be the main initial revenue stream.

For 2G systems, the Average Revenue per User (ARPU) has dropped (and is dropping) rapidly as the market saturates and competition bites. For example, Analysis predicts that the European ARPU will continue to decline, halving over the next 10 years from about 30 Euros per month in 2001. They also suggest that a 3G operator cannot make a satisfactory return on voice alone, because their cumulative cash flow only becomes positive in 2010.

If an operator cannot be profitable from voice alone, it clearly must increase the revenue considerably with additional services. Since these are likely to be data services of one form or another, the extra revenue required is often called the 'data gap'. Many services have been suggested to bridge this 'data gap', which will be discussed shortly.

Lessons from 2.5G - I-mode, WAP and GPRS

The data capability enhancements that have been added on to 2G systems can be viewed as a stepping stone to 3G - and hence they are collectively called '2.5G': an intermediate point in terms of technology (bit rates, etc.) and commerce (the chance to try out new services, etc.).

Undoubtedly, the most successful so far has been i-mode in Japan. I-mode allows users to do their e-mail and text messaging. Other popular activities include viewing news and horoscopes, and downloading ring tones, cartoon characters and train times. Users can connect to any site written in cHTML (compact HTML - a subset of HTML (Hyper-Text Markup Language) designed so that pages can display quickly on the small screens of the I-mode terminals), but some sites are approved by NTT-DoCoMo (the operator); these have to go through a rigorous approval process, e.g. content must be changed very regularly. The belief is that if users can be confident that sites are 'good', that will encourage extra traffic and new subscribers in a virtuous circle for the operators, content providers and customers. Current download speeds are limited to 9.6 kbit/s with an upgrade to 28.8 kbit/s planned for spring 2002.

I-mode has grown very rapidly from its launch in February 1999 to over 28 million users in October 2001. The basic charge for i-mode is about 300 Yen (\$2.50) per month, plus 2.4 Yen (2 cents) per k-byte downloaded. The DoCoMo-approved 'partner sites' have a

further subscription charge of up to about 300 Yen (\$2.50) per month, which is collected via the phone bill, with DoCoMo retaining 9% as commission . For other sites, DoCoMo just receives the transport revenues.

GSM's WAP (Wireless Application Protocol) is roughly equivalent to i-mode, but has been far less successful, with fewer than 10% of subscribers. The Economist suggests various reasons for i-mode's relative (and absolute) success, for example:

- Low PC penetration in Japan (for cultural reasons).
- High charges for PSTN dial-up access in Japan.
- The Japanese enthusiasm for gadgets.
- Non-standardization of i-mode - Meaning that an operator can launch a new service more easily, including specifying to manufacturers what handsets they want built (e.g. with larger LCD screens).
- Expectation management - This was sold to users as a special service (with applications and content useful for people 'on the move'), whereas WAP was (over) hyped as being 'just like the Internet'.
- Its business model - This provides a way for content producers to charge consumers.

GPRS, which is a packet data service being added on to GSM networks, has started rolling out during 2001. It will eventually offer connections at up to 144 kbit/s, but 14–56 kbit/s to start with. Like i-mode, GPRS is an 'always on' service. Again, this is likely to provide important lessons as to what sort of services are popular with consumers and businesses, and how to make money out of them.

3G Services

Many services have been suggested for 3G in order to bridge the 'data gap' discussed earlier, and so provide sufficient revenue to more than cover the costs outlined above. Typical services proposed are m-commerce, location-based services and multimedia (the integration of music, video, and voice such as video-phones, video-on-demand and multimedia messaging).. It is generally accepted that a wide range of services is required there is no single winner - but there are different views as to which will prove more important than others. For example:

- **Multimedia Messaging** - Text messaging (e.g. SMS) has been very successful, and on the Internet we are seeing a rapid growth in 'instant messaging' (IM) - for example, AOL's Instant Messenger and ICQ services each have over 100 million registered users . In particular, it is predicted that the multimedia messaging service (MMS) will become very popular in 3G. For example, Alatto believe that the primary data revenue source will be MMS. Typical MMS applications might be the sharing of video clips and music; similar ideas have proved very already popular on the Internet, e.g. Napster. 3G terminals are likely to include a camera and appropriate display exactly to enable services like these. In a similar vein, but using wireless LAN technology instead of 3G, Cybiko includes MMS to nearby friends. (Cybiko is a wireless hand-held computer for teens.)

- **Location-based services** - An operator knows the location of a mobile user, and thus services can be tailored to them. For example, 'where is the nearest Thai restaurant?'; the reply can include a map to guide you there and an assurance that a table is free. Early examples are available today, for instance J-phone's J-Navi service. Analysis expects that 50% of all subscribers will use such services, with a global revenue of \$18.5 billion by the end of 2006.
- **M-commerce** - This is e-commerce to mobile terminals, for example, ordering goods or checking your bank account. Durlacher predicted the European m-commerce market to grow from Euro 323 million in 1998 to Euro 23 billion by 2003. Sonera have trialed a service where drinks can be bought from a vending machine via a premium-rate GSM phone number or SMS message. M-commerce will grow as techniques for collecting micro-payments are developed and refined. One possible option is to have these collected by your service provider and added and billed using either pre- or post-pay. Smart cards, including SIM cards, could be used to authenticate these transactions. Another m-commerce application is personalized advertising, i.e. tailored to the user.
- **Business-to-business m-commerce** - This will allow staff working at a customer's site to obtain information from their company's central database, to provide quotes and confirm orders on the spot. This could help to cut their costs (less infrastructure and fewer staff whom it is easier to manage) as well as provide a better service to the customer

As well as the extra revenue from these new services, operators hope that they will encourage customers to make more voice calls and also that by offering different, innovative services, they will reduce customer 'churn' - i.e. customers will be more likely to stick with them. Such an impact does seem to have happened with I-mode.

Overall Business Case for 3G

The reason that there is so much interest in 3G and the mobile Internet is summarized very well by Stand age: The biggest gamble in business history; control of a vast new medium; the opportunity at last to monetize the Internet: clearly, a great deal is at stake. Some say it is all just wishful thinking. But in many parts of the world not only Japan millions of people are even now using phones and other handheld devices to communicate on the move. All over the globe, the foundations for this shift to more advanced services are already in place.

Here, we are not interested in developing the business case *per se* - only to show that any technology that improves the business case must be a good thing and to point out the areas where we believe IP technologies can make a difference.

3G Value Chain

A value chain is a map of the companies involved in delivering services to the end consumer and is drawn up to identify who makes the profits (in business-speak, making a profit is called 'value generation').

Lessons from 2G

The 2G value chain is pretty simple - basically, users buy handsets and billing packages from operators through retail outlets. The importance of terminal manufacturers has been strengthened by operators subsidizing handsets, "effectively supporting terminal manufacturers' brands (e.g. Nokia) to the extent that these now outweigh the brands of the operator in customers' minds". The content - voice and SMS - is generated by the users themselves. Recently, a slight addition to the chain has been 'virtual operators'; this is basically about branding, and means that (taking a UK example) a user buys a Virgin phone that is actually run by One 2 One (the real operator).

In 2G, the operators control the value chain and the services offered via the SIM card. This is sometimes called the 'walled garden' approach - the operator decides what flowers (services) are planted in the garden (network) and stops users seeing flowers in other gardens the other side of the wall

Possible 3G Value Chain

For 3G networks, it is often suggested that the value chain will become more complicated. Many possibilities have been suggested, and [Figure 1.2](#) shows one possibility by Harmer and Friel. They suggest that the roles of the players are as follows:

- Network operator - Owns the radio spectrum and runs the network.
- Service provider - Buys wholesale airtime from the network operator and issues SIM cards and bills.
- Mobile Virtual Network Operator (MVNO) - MVNOs own more infrastructure than service providers - perhaps some switching or routing capacity.
- Mobile Internet Service Provider (M-ISP) - Provide users with IP addresses and access to wider IP networks.
- Portal Provider - Provide a 'homepage' and hence access to a range of services that are in association with the portal provider.
- Application Provider - Supplies products (e.g. software) that are downloaded or used on line.
- Content provider - Owners of music or web pages and so forth.



Figure 1.2: Possible 3G value chain. Source: Harmer & Friel

Of course, there are many other possible models (see, for example), and it must also be pointed out that some of these 'logically' different roles might actually be played by the same operator. Indeed, it is not unrealistic to think that many 3G operators - those owning licenses could play all the roles (except, of course, that of MVNO).

Some people believe that the value will shift, compared with 2G, from network operators to content providers, especially following the success of i-mode. For example, KPMG estimate that "only 25% of the total revenue will be in the transmission of traffic and the remaining 75% will be divided up among content creation, aggregation, service provision, and advertising" . However, there is disagreement about who in the value chain will benefit:

1.4.2 Impact of 'IP for 3G' on Business Case

The key impact that 'IP for 3G' could have is to help the convergence of the Internet and communications. Cleevly speculates that it could lead to a fall in the unit cost of communications by a factor of nearly 1000 by 2015, because convergence will cause a massive growth in demand and hence large economies of scale. The following gives some 3G perspective.

Costs

IP is becoming the ubiquitous protocol for fixed networks, so economies of scale mean that it is very likely that IP-based equipment will be the cheapest to manufacture and buy for mobile networks. Further, an operator that runs both fixed and mobile network services should be able to roll out a single, unified network for both jobs, leading to savings on capital costs and maintenance. It should also allow the reuse of standard Internet functionality for things like security. IP evolution in both fixed and mobile networks offers the possibility of having a single infrastructure for all multimedia delivery - to any terminal over any access technology. This will not necessarily drive down costs for any one particular service: after all, the PSTN is supremely optimized for voice delivery, but for future multimedia services where voice, video, real-time, non-real-time and multicast all mix together, IP evolution of both the fixed and mobile networks to a common architecture holds out the prospect of lower costs.

Services and Revenues

From an end user's perspective, applications are increasingly IP-based. In an all-IP network, the same applications will be available for mobile users as for fixed, and they will behave as intended. Existing applications will not need to be rewritten for the special features of the mobile system (as tends to happen today). Another issue is security, which is critical for M-commerce applications. 'Mobile specials' may lead to new security holes that need plugging as they become apparent, and also users have to be re-convinced that their e-commerce transactions are secure. WAP provides an example of this problem.

The Internet is adding call/session control, particularly via the Session Initiation Protocol (SIP). As well as enabling peer-to-peer calls, which are certainly needed in 3G, this elegant and powerful protocol will enable service control similar to that of the 'intelligent network': things like 'ring back when free' and other supplementary services, or more complex things like 'divert calls from boss to answer phone whilst I am watching cricket on Internet-TV'. Again, an 'IP for 3G' approach should mean that the user experience is the same regardless of whether they are on a fixed or mobile network. More speculatively, 'IP for 3G' might enable the same location-based services to be offered more easily on the *fixed* network as well.

Overall, 'IP for 3G' should mean that new applications can concentrate on the particular benefits of mobility, such as location-based services. This will give benefits for the user (obtaining the applications that the user desires and is familiar with) and for the application writer (lower development costs, wider market - and hence a wider choice of applications for the user). Hence, companies gain the extra traffic and extra revenues they want.

Value Chain

The impact of IP on the 3G value chain is unclear. There is some tension between the 2G walled garden approach and that of the Internet where anyone can set up a web server and deliver services to whoever discovers it. i-mode is an interesting half-way house, with its partner sites, but also allowing access to any site. Further, the Internet approach allows services to run over any link layer (bit transport mechanism), whereas 3G's stovepipe approach clearly locks the user into the 3G air interface. The impact of other high-speed wireless technologies (such as wireless LANs, Blue-tooth, and a future system using a re-farmed analogue TV spectrum) is very interesting and uncertain. It is not at all obvious whether they should be viewed as a threat to 3G (they take traffic away from the user), or as a complement (they enhance the capacity and coverage), or even as a benefit (they get people hooked on the 3G services, which is what they make money on).

1.5 Conclusion

In this chapter, we started by outlining fairly broad definitions of 'IP' and by '3G':

- 'IP' is about the Internet, its design principles, protocols and standardization approach.
- '3G' is about the new mobile system, its architecture, network, and air interface.

So, 'IP for 3G' is about the convergence of the Internet and mobile communications revolutions. This report concentrates on technological, and especially network, aspects of this convergence. The first chapter has given some motivation for why we believe that IP for 3G is important. The reasons fall into two categories:

- Engineering - Essentially about why IP's design principles are a good thing, focusing on IP's clear protocol layering and the end-to-end principle.
- Economic - About how IP can dramatically reduce the costs of building the mobile multimedia network - from the benefits of integration and economies of scale and can increase the range of services it carries.

The two sets of reasons are closely connected - it is IP's good engineering design principles that enable the network to be much cheaper and the services offered on it far more numerous. We believe that the flexibility of an all-IP mobile network will liberate application developers from having to understand the details of the network, so that they can concentrate on what the end users want - indeed, there is the flexibility just to try ideas out until they haphazardly discover things that people like. This process will ignite a Cambrian explosion of applications and services. It will lead to a dramatic increase in users and traffic which in turn will lead to further economies of scale and cost reductions

So, 'IP for 3G' is in effect our campaign slogan - we believe that there should be *more IP in 3G*. However, adding IP technologies and protocols into 3G is not trivial - there are many difficulties and unresolved issues. So, 'IP for 3G' is an interesting and important topic that requires further study and research.

Chapter 2

Introduction to 3G Networks

2.1 Introduction

What exactly are 3G networks? 3G is short for Third Generation (Mobile System). Here is a quick run-down:

- 1G, or first generation systems, were analogue and offered only a voice service each country used a different system, in the UK TACS (Total Access Communications System) was introduced in 1980. 1G-systems were not spectrally efficient, were very insecure against eavesdroppers, and offered no roaming possibilities (no use on holidays abroad.).
- 2G heralded a digital voice and messaging service, offered encrypted transmissions, and was more spectrally efficient than 1G. GSM (Global System for Mobile communication) has become the dominant 2G standard and roaming is now possible between 150+ countries where GSM is deployed.
- 3G - if the popular press is to be believed - will offer true broadband data: video on demand, videophones, and high bandwidth games will all be available soon. 3G systems differ from the second generation voice and text messaging services that everybody is familiar with in terms of both the bandwidth and data capabilities that they will offer. 3G systems are due to be rolled out across the globe between 2002 and 2006. 3G will use a new spectrum around 2 GHz, and the licenses to operate 3G services in this spectrum have recently hit the headlines because of the huge amounts of money paid for licenses by operators in the UK and Germany (£50 billion or so). Other countries have raised less or given away licenses in so-called 'beauty contests' of potential operators.

3G systems might be defined by: the type of air interface, the spectrum used, the bandwidths that the user sees, or the services offered. All have been used as 3G definitions at some point in time. In the first wave of deployment, there will be only two flavors of 3G - known as UMTS (developed and promoted by Europe and Japan) and cdma2000 (developed and promoted by North America). Both are tightly integrated systems that specify the entire system - from the air interface to the services offered. Although each has a different air interface and network design, they will offer users broadly the same services of voice, video, and fast Internet access.

3G (and indeed existing second generation systems such as GSM) systems can be divided very crudely into three (network) parts: the air interface, the radio access network, and the core network. The air interface is the technology of the radio hop from the terminal to the base station. The core network links the switches/routers together and extends to a gateway linking to the wider Internet or public fixed telephone network. The Radio Access Network (RAN) is the 'glue' that links the core network to the base stations and deals with most of the consequences of the terminal's mobility.

This chapter concerns the core and access networks of 3G systems - because that is where IP (a network protocol) could make a difference to the performance and architecture of a 3G network. The chapter first reviews the history of 3G developments from their 'conception' in the late 1980s, through their birth in the late 1990s, to the teething troubles that they are currently experiencing. The history of 3G development shows that the concepts of 3G evolved significantly as the responsibility for its development moved from research to standardization - shedding light on why 3G systems are designed the way they are. Included in this section is also a 'who's who' of the standards world, a very large number of groups, agencies, and for a have been, and still are, involved in the mobile industry. In the second half of the chapter, we introduce the architecture of UMTS (the European/Japanese 3G system) and look at how the main functional components QoS, mobility management, security, transport and network management are provided. A short section on the US cdma2000 3G system is also included at the end of the chapter.

The purpose of this chapter is to highlight the way UMTS (as an example 3G systems) works at a network level - in terms of mobility management, call control, security, and so forth. This is intended as a contrast with the descriptions of how IP research is evolving to tackle these functions in the chapters that follow. The final chapter combines the two halves - IP and 3G - to pursue the main argument of the report - that 3G should adopt IP design principles, architectures and protocols thereby allowing greater efficiency, fixed mobile convergence, and new IP services (e.g. multicast).

2.2 Mobile Standards.

Mobile systems development, particularly that of 3G systems, is inextricably bound up with the process of standardization,. Why? Why is standardization so important? The best answer to that question is probably to look at GSM, whose success could reasonably be described as the reason for the vast interest and sums of money related to 3G. GSM was conceived in the mid-1980s - just as the first analogue cellular mobile systems were being marketed. These analogue systems were expensive and insecure (easy to tap), and there was no inter-working between the great variety of different systems (referred to as 'first generation systems') deployed around the world. GSM introduced digital transmission that was secure and made more efficient use of the available spectrum. What GSM offered was a tight standard that allowed great economies of scale and competitive procurement. Operators were able to source base stations, handsets, and network equipment from a variety of suppliers, and handsets could be used anywhere the GSM standard was adopted. The price of handsets and transmission equipment fell much faster than general trends in the electronics industry. GSM also offered a roaming capability since the handsets could be used on any GSM system; made possible by a remote authentication facility to the home network. There were other advantages of moving to a digital service, such as a greater spectral efficiency and security, but in the end, it was the mass-market low cost (pre-pay packages have sold for as little as £20) that was the great triumph of GSM standardization. In terms of world markets, GSM now accounts for over 60% of all second generation systems and has 600 million users in 150 countries; no other system has more than 12%

However, the standardization process has taken a very long time, 18 years from conception (1980) to significant penetration (say 1998). It has resulted in a system that is

highly optimized and integrated for delivering mobile voice services and is somewhat difficult to upgrade. As an example, consider e-mail: e-mail has been in popular use since, maybe, 1992 but 10 years on, how many people can receive e-mail on their mobile? This facility is beginning to appear - along with very limited web-style browsing on mobiles [e.g. using WAP (Wireless Application Protocol) and i-mode in Japan]. Standards can also be a victim of their own success - 2G (and GSM in particular) has been so successful that operators and manufacturers have been keen to capitalize on past investments and adopt an evolutionary approach to the 3G core network.

2.2.1 Who's who in 3G Standards

At this point, it is perhaps a good idea to provide a brief 'who's who' to explain recent developments in the standards arena.

- 3GPP - In December 1998, a group of five standards development organizations agreed to create the Third Generation Partnership Project (3GPP - www.3gpp.org). These partners were: ETSI (EU), ANSI-TT (US), ARIB and TTC (Japan), TTA (Korea), and CWTS (China). Basically, this was the group of organizations backing UMTS and, since August 2000, when ETSI SMG was dissolved, has been responsible for all standards work on UMTS. 3GPP have now completed the standardization of the first release of the UMTS standards - Release 99 or R3. GSM upgrades have always been known by the year of standardization, and UMTS began to follow that trend, until the Release 2000 got so behind schedule that it was broken into two parts and renamed R4 and R5. In this chapter, only the completed R3 (formally known as Release 99) will be described. Chapter 6 looks at developments that R4 and R5 will bring. 3GPP standards can be found on the 3GPP website - www.3GPP.org - and now completely specify the components and the interfaces between them that constitute a UMTS system.
- 3GPP2 - 3GPP2 (www.3gpp2.org) is the cdma2000 equivalent of 3GPP with ARIB and TTC (Japan), TR45 (US), and TTA (Korea). It is currently standardizing cdma2000 based on evolution from the cdmaOne system and using an evolved US DAMPS network core. (The latter part of this chapter gives an account of packet transfer in cdma2000.)
- ITU - The International Telecommunications Union (ITU - www.itu.int) was the originating force behind 3G with the FLMTS concept (pronounced Flumps and short for Future Land Mobile Telecommunication System) and work towards spectrum allocations for 3G at the World Radio Conferences. The ITU also attempted to harmonize the 3GPP and 3GPP2 concepts, and this work has resulted in these being much more closely aligned at the air interface level. Currently, the ITU is just beginning to develop the concepts and spectrum requirements of 4G, a subject that is discussed at length in [Chapter 6](#).
- IETF - The Internet Engineering Task Force (www.ietf.org) is a rather different type of standards organization. The IETF does not specify whole architectural systems, rather individual protocols to be used as part of communications systems. IETF protocols such as SIP (Session Initiation

Protocol) and header compression protocols have been incorporated in to the 3GPP standards. IETF meetings take place three times a year and are completely open, very large (2000+ delegates), and very argumentative (compared with the ITU meeting, say). Anyone can submit an Internet draft to one of the working groups, and this is then open to comments. If it is adopted, it becomes a Request For Comments (RFC); if not, it is not considered any further.

- MWIF - The industry pressure group Mobile Wireless Internet Forum (www.mwif.org) comprises operators, manufacturers, ISPs (Internet Service Providers) and Internet equipment suppliers. MWIF, since early 2000, has been producing a functional architecture that separates the various components of a 3G systems - for example, the access technology to provide opportunities for IP technologies such as Wireless LANs to be used.

2.3 History of 3G

It is not widely known that 3G was conceived in 1986 by the ITU (International Telephony Union). It is quite illuminating to trace the development of the ideas and concepts relating to 3G from conception to birth. What is particularly interesting, perhaps, is how the ideas have changed as they have passed through different industry and standardization bodies. 3G was originally conceived as being a single world-wide standard and was originally called FLMTS (pronounced Flumps and short for Future Land Mobile Telecommunication System) by the ITU. By the time it was born, it was quin's five standards and the whole project was termed the IMT-2000 family of standards. After the ITU phase ended in about 1998, two bodies 3GPP and 3GPP2 completed the standardization of the two flavors of 3G that are actually being deployed today and over the next few years (UMTS and cdma2000, respectively). Meanwhile, these bodies, along with the Operator Harmonization Group (OHG), are looking at unifying these into a single 3G standard that allows different air interfaces and networks to be 'mixed and matched'.

It is convenient to divide up the 3G gestation into three stages (trimesters):

- Pre-1996 - The Research Trimester.
- 1996–1998 - The IMT-2000 Trimester.
- Post-1998 - The Standardization Trimester.

2.3.1 Pre-1996 - The Research Trimester

Probably the best description of original concept of 3G can be found in Alan Clapton's quote head of BT's 3G development at the time.

"3G ...The evolution of mobile communications towards the goal of universal personal communications, a range of services that can be anticipated being introduced early in the next century to provide customers with wireless access to the information super highway and meeting the 'Martini' vision of communications with anyone, anywhere and in any medium."

Here are the major elements that were required to enable that vision:

- A world-wide standard - At that time, the European initiative was intended to be merged with US and Japanese contributions to produce a single world-wide system known by the ITU as FLMTS. The vision was a single hand-set capable of roaming from Europe to America to Japan.
- A complete replacement for all existing mobile systems - UMTS was intended to replace all second generation standards, integrate cordless technologies as well as satellite (see below) and also to provide convergence with fixed networks.
- Personal mobility - Not only was 3G to replace existing mobile systems, but its ambition stretched to incorporating fixed networks as well. Back in 1996, of course, fixed networks meant voice, and it was predicted in a European Green Paper on Mobile Communications that mobile would quickly eclipse fixed lines for voice communication. People talked of Fixed Mobile Convergence (FMC) with 3G providing a single bill, a single number, common operating, and call control procedures. Closely related to this was the concept of the Virtual Home Environment (VHE).
- Virtual Home Environment - The virtual home environment was where users of 3G would store their preferences and data. When a user connected, be it by mobile or fixed or satellite terminal, they were connected to their VHE, which then was able to tailor the service to the connection and terminal being used. Before a user was contacted, the VHE was interrogated, so that the most appropriate terminal could be used, and the communication tailored to the terminals and connections of the parties.
- Broadband service (2 Mbit/s) with on-demand bandwidth - Back in the early 1990s, it was envisaged that 3G would also need to offer broadband services – typically meaning video and video telephony. This broadband requirement meant that 3G would require a new air interface, and this was always described as broadband and typically thought to be 2 Mbit/s. Associated with this air interface was the concept of bandwidth on demand - meaning that it could be changed during a call. Bandwidth on demand could be used, say, to download a file during a voice conversation or upgrade to a higher-quality speech channel mid-way through a call.
- A network based on B-ISDN - Back in the early 1990s, another concept certainly at BT, was that every home and business would be connected directly to a fiber optic network. ATM transport and B-ISDN control would then be used to deliver broadcast and video services, an example being video on demand whereby customers would select a movie, and it would be transmitted directly to their home. B-ISDN [Broadband ISDN was supposed to be the signaling for a new broadband ISDN service based on ATM transport - it was never actually developed, and ATM signaling is still not yet sufficiently advanced to switch circuits in real time. ATM (asynchronous transfer mode) is explained in the latter part of this chapter: it is used in the UMTS radio access and core networks.] Not

surprisingly, given the last point, it was assumed that the 3G network would be based on ATM/B-ISDN.

- A satellite component - 3G was always intended to have an integrated satellite component, to provide true world-wide coverage and fill in gaps in the cellular networks. A single satellite/3G handset was sometimes envisaged. (Surprisingly, since satellite handsets tend to be large).

The classic picture seemingly compulsory in any description of 3G is of a layered architecture of radio cells (Figure 2.1). There are mega-cells for satellites, macro-cells for wide area coverage (rural areas), micro-cells for urban coverage, and Pico-cells for indoor use. There is a mixture of public and private use and always a satellite hovering somewhere in the background.



Figure 2.1: Classic 3G layer diagram.

In terms of forming this vision of 3G, much of the early work was done in the research programs of the European Community, such as the RACE (Research and development in Advanced Communications technologies in Europe) programs with projects such as MONET (looking at the transport and signaling technologies for 3G) and FRAMES (evaluating the candidate air interface technologies). In terms of standards, ETSI (European Telecommunications Standards Institute) completed development of GSM phase 2, and at the time, this was intended to be the final version of GSM and for 3G to totally supersede it and all other 2G systems. As a result, European standardization work on 3G, prior to 1996, was carried out within an ETSI GSM group called, interestingly, SMG5 (Special Mobile Group).

2.3.2 1996–1998 - The IMT 2000 Trimester

It is now appropriate to talk of UMTS (Universal Mobile Telecommunications System) as the developing European concept was being called. In the case of UMTS, the Global Multimedia Mobility report was endorsed by ETSI and set out the framework for UMTS standardization. The UMTS Forum - a pressure group of manufacturers and operators produced the influential UMTS forum report (www.umts-forum.org) covering all non standardization aspects in UMTS such as regulation, market needs and spectrum requirements. As far as UMTS standardization was concerned, ETSI transferred the standardization work from SMG5 to the various GSM groups working on the air interface, access radio network, and core network.

In Europe, there were five different proposals for the air interface most easily classified by their Medium Access Control (MAC) schemes, in other words, how they allowed a

number of users to share the same spectrum. Basically, there were time division (TDMA Time Division Multiple Access), frequency division (OFDM - Orthogonal Frequency Division Multiple Access), and code division proposals (CDMA). In January 1998, ETSI chose two variants of CDMA - Wideband CDMA (W-CDMA) and time division (TD-CDMA) the latter basically a hybrid with both time and code being used to separate users. W-CDMA was designated to operate in paired spectrum [a band of spectrum for up link and another (separated) band for down link] and is referred to as the FDD (Frequency Division Duplex) mode, since frequency is used to differentiate between the up and down traffic. In the unpaired spectrum, a single monolithic block of spectrum, the TD-CDMA scheme was designated, and this has to use time slots to differentiate between up and down traffic (FDD will not work for unpaired spectrum - see Section 2.4 for more details), and so is called the TDD (Time Division Duplex) mode of UMTS.

In comparison, GSM is a FDD/TDMA system - frequency is used to separate up and down link traffic, and time division is used to separate the different mobiles using the same up (or down) frequency.

Part of the reason behind the decision to go with W-CDMA for UMTS was to allow harmonization with Japanese standardization.

Unfortunately, in North America, the situation was more complicated; firstly, parts of the 3G designated spectrum had been licensed to 2G operators and other parts used by satellites; secondly, the US already has an existing CDMA system called cdmaOne that is used for voice. It was felt that a CDMA system for North America needed to be developed from cdmaOne with a bit rate that was a multiple of the cdmaOne rate. Consequently, the ITU recognized a third CDMA system - in addition to the two European systems - called cdma2000. It was also felt that the lack of 3G spectrum necessitated an upgrade route for 2G TDMA systems resulting in a new TDMA standard called UMC-136, which is effectively identical to a proposed enhancement to GSM called EDGE (Enhanced Data rates for Global Evolution). This takes advantage of the fact that the signal-to-noise ratio (and hence potential data capacity) of a TDMA link falls as the mobile moves away from the base station. Users close to base stations essentially have such a good link that they can increase their bit rate without incurring errors. By using smaller cells or adapting the rate to the signal-to-noise ratio, on average, the bit rate can be increased. In CDMA systems, the signal-to-noise ratio is similar throughout the cell.

Finally the DECT (Digital European Cordless Telecommunications) - developed by ETSI for digital cordless applications and used in household cordless phones, for example - inhabits the 3G spectrum and has been included as the fifth member of the IMT-2000 family of 3G standards (Table 2.1) as the ITU now called the FPLMTS vision.

Table 2.1: IMT 2000 family of 3G standards		
IMT2000 designation	Common term	Duplex type
IMT-DS Direct Sequence CDMA	Wideband CDMA	FDD
IMT-MC Multi Carrier CDMA	Cdma2000	FDD
IMT-TD Time Division CDMA	TD/CDMA	TDD
IMT-SC Single Carrier	UMC-136 (EDGE)	FDD
IMT-FS Frequency Time	DECT	TDD

During this period, 3G progressed from its 'Martini' vision, 'anytime, anyplace, anywhere', to a system much closer, in many respects, to the existing 2G networks. It is true that the air interface was a radical change from TDMA - it promised a better spectral efficiency, bandwidth on demand, and broadband connections - but the core networks chosen for both UMTS and cdma2000 were based on existing 2G networks: in the case of UMTS, an evolved GSM core, and for cdma2000, an evolved ANSI-41 core (another time division circuit switching technology standard). The major reason for this was the desire by the existing 2G operators and manufacturers to reuse as much existing equipment, development effort, and services as possible. Another reason was the requirement for GSM to UMTS handover, recognizing that UMTS coverage will be limited in the early years of roll-out.

The radio access network for UMTS was also new, supporting certain technical requirements of the new CDMA technology and also the resource management for multimedia sessions. The choice of evolved core network for UMTS is probably the key non-IP friendly decision that was taken at this time, meaning that that UMTS now supports both IP and X25 packets using a common way of wrapping them up and transporting them over an underlying IP network. (X25 is an archaic and heavyweight packet switching technology that pre-dates IP and ATM). In the meantime, X25 has become totally defunct as a packet switching technology, and IP has become ubiquitous, meaning that IP packets are wrapped up and carried within outer IP packets because of a no-longer useful legacy requirement to support X25.

2.3.3 1998 Onwards - The Standardization Trimester

After 1998, the function of developing and finalizing the standards for UMTS and cdma2000 passed to two new standards bodies: 3GPP and 3GPP2, respectively. These bodies have now completed the first version (or release) of the respective standards (e.g. R3 - formally known as Release 99 for UMTS), and these are the standards that equipment is currently being procured against for the systems currently on order around the world. Current order numbers are UMTS 34, cdma2000 9, and EDGE 1 (number of systems)

2G systems have not stood still and are introducing higher-speed packet data services (so called 2.5G systems: the GSM 2.5 G evolution is GPRS - GSM Packet Radio System). These will offer either subscription or per-packet billing and allow users to be 'always on' without paying a per-second charge as they currently do for circuit-based data transfer. The new network elements needed to add packet data to GSM are also needed for UMTS, and details of these are given later in the chapter (for a good description of GPRS,

In early 2000, 3G license auctions raised £50 billion in the UK and Germany, and many expected that services would be universally available by 2002. That now looks unlikely with the major downturn in the telecoms industry, the failure of WAP to take off in Europe, and technical delays over the new air interfaces and terminals. After WAP was widely rejected because of long connection times and software errors, many operators are using 2.5G systems - such as GPRS - as a proving ground for 3G. NTT launched a limited 3G service in Tokyo, in late 2001, with a few hundred handsets. Most commentators now see 3G deployment held back until 2004 and much site and infrastructure sharing to produce cost savings. Since the first UMTS Release, there has been work in groups like 3GIP to be more revolutionary and include more IP (in its widest sense) in 3G. 3GIP has produced a number of technical inputs to the second version of UMTS - originally called Release 2000 but now broken into two releases, known as R4 and R5 in the revised (so as to avoid the embarrassment of finishing Release 2000 in 2002) numbering scheme. We shall look at what R4 and R5 offer in Chapter 6.

Finally the operator harmonization group and 3GPP/3GPP2 are working to harmonize UMTS, cdma2000, and EDGE such that any of these air interfaces and their associated access networks - or indeed a Wireless LAN network - can be connected to either an IS-41 or evolved GSM core network. The final goal is a single specification for a global 3G standard.

2.4 Spectrum - The 'Fuel' of Mobile Systems

Now is a good time to consider spectrum allocation decisions, as these have a key impact on the 3G vision in terms of the services (e.g. bandwidth or quality) that can be provided and the economics of providing them.

In any cellular system, a single transmitter can only cover a finite area before the signal-to noise ratio between the mobiles and base stations becomes too poor for reliable transmission. Neighboring base stations must then be set up and the whole area divided into cells on the basis of radio transmission characteristics and traffic density. The neighboring cells must operate on a different frequency (e.g. GSM /D-AMPS) or different spreading code (e.g. WCDMA or cdma One; see [Figure 2.2](#)). Calls are handed over between cells by arranging for the mobile to use a new frequency, code or time slot. It is a great, but profitable and very serious, game of simulation and measurement to estimate and optimize the capacity of different transmission technologies. For example, it was originally estimated that W-CDMA would offer a 10-fold improvement in transmission efficiency (in terms of bits transmitted per Hertz of spectrum) over TDMA (Time Division Multiple Access - such as GSM and D-AMPS) in practice, this looks to be twofold at best.




Figure 2.2: Typical (TDMA) cellular system.

Typically, figures for a 2G system are:

- Bandwidth of a call - 14 kbit/s (voice).
- Bandwidth available 30 MHz (Orange - UK).
- Efficiency 0.05 (or frequency reuse factor of 20 - meaning that one in 20 cells can use the same frequency with acceptable interference levels).

Any capacity can be achieved by simply building a higher base station density (although this increases the costs). Second, the higher the bandwidth per call, the lower the capacity so broadband systems offering 2 Mbit/s to each user need about 150 times the spectrum bandwidth of voice systems to support the same number of users (or will support around 150 times less users), all other things being equal. Third, any major increase in efficiency for a given capacity means that either a smaller density of base stations or fewer spectrums is required, and, given both are very expensive, this is an important research area. Unfortunately for 3G systems, as mentioned above, this factor has improved by only 2 over current GSM systems. Finally if the bandwidth of a voice call can be halved, the capacity of the system can be doubled; this is the basis of introducing half-rate (7 kbit/s) voice coding in GSM.

So, given this analysis, it is hard to escape the conclusion that 3G systems need a lot of spectrum. However, radio spectrum is a scarce resource. To operate a cellular mobile system only certain frequencies are feasible: at higher frequencies, radio propagation characteristics mean that the cells become smaller, and costs rise. For example, 900-MHz GSM operators (e.g. Cellnet in the UK) require about half the density of stations - in rural areas compared with 1800-MHz GSM operators like Orange. Also, above about 3 GHz, silicon technology can no longer be used for the transmitters and receivers necessitating a shift to gallium arsenide technology, which would be considerably more expensive. The difficulties of finding new spectrum in the 500-3000-MHz range should not be under-emphasized see for a lengthy account of the minutiae involved - but, in short, all sorts of military, satellite, private radio and navigation systems, and so forth all occupy different parts of the spectrum in different countries. Making progress to reclaim or “re-farm” as it is known the spectrum is painfully slow on a global scale. The spectrum bands earmarked for FPLMTS at the World Radio Conference in 1992 were 1885–2025 MHz and 2110–2200 MHz, a total of 230 MHz. However, a number of factors and spectrum management decisions have since eroded this allocation in practice:

- Mobile satellite bands consume 2 x 30 MHz.
- In the US, licenses for much of the FPLMTS band have already been sold off for 2G systems.

- Part of the bands (1885–1900 MHz) overlap with the European DECT system.

All of this means that only 2×60 MHz and an odd 15 MHz of unpaired spectrum are available for 3G in Europe and much less in the US. The paired spectrum is important this means equal chunks of spectrum separated by a gap - one part being used for up link communications and the other for down link transmission. Without the gap separating them up and down link transmissions would interfere at the base station and mobile if they transmitted and received simultaneously. By comparison, in the UK today, 2×100 MHz is available for GSM, shared by four operators. Figure 2.3 shows the general world position on the 3G spectrum - explaining why many commentators expect 3G to be much less influential in the US and rolled out earlier in Europe and Japan.




Figure 2.3: Global spectrum allocations for 3G (MSS bands are satellite spectrum).

In the UK auction/licensing process, there were a dozen or so bidders chasing five licenses, resulting in three getting 10 MHz and two buying 15 MHz of paired spectrum per operator. BT has acquired 2×10 MHz of paired spectrum and 5 MHz of unpaired spectrum. BT Cellnet will use the paired spectrum with 5 MHz for macro cells and 5 MHz for micro cells there being no need for frequency planning in a W-CDMA system.

Chapter 3

Universal Mobile Telecommunications System (UMTS)

Network Overview

In order to illustrate the operation of a UMTS network, this section describes a day in the life of a typical UMTS user - this sort of illustration is often called a usage case or a scenario. The major network elements - the base stations and switches etc. - will be introduced, as well as the functionality that they provide. This at least has the merit of avoiding a very sterile list of the network elements and serves as a high-level guide to the detailed description of UMTS functionality that follows.

Mary Jones is 19 years old and has just arrived at the technical Polytechnic of Darmstadt. She is lucky that her doting father has decided to equip her with a 3G terminal before allowing her to live away from home, but then this is 2004, and such terminals are now common in Germany and much of Europe.

Mary first turns her terminal on after breakfast and is asked to enter her personal PIN code. This actually authenticates her to the USIM (UMTS Subscriber Identity Module) a smart card that is present within her terminal. The terminal then searches for a network, obtains synchronization with a local base station, and, after listening to the information on the cell's broadcast channel, attempts to attach to the network. Mary's subscription to T-Nova is based on a 15-digit number (which is not her telephone number) identifying the USIM inside her terminal. This number is sent by the network to a large database called the home location register (HLR) located in the T-Nova core network. Both the HLR and Mary's USIM share a 128-bit secret key this is applied by the HLR to a random number using a one-way mathematical function (one that is easy to compute but very hard to invert). The result and the random number are sent to the network, which challenges Mary's USIM with the random number and accepts her only if it replies with the same result as that sent from the HLR (Figure 3.1)




Figure 3.1: UMTS Architecture.

After attaching to the network, Mary decides to call her dad - perhaps, although unlikely, to thank him for the 3G terminal. The UMTS core network is divided into two halves one half dealing with circuit-switched (constant bit rate) calls - called the circuit-switched domain and the other - the packet-switched domain - routing packets sessions. At this time, Mary attempts to make a voice call, and her terminal utilizes the connection management functions of UMTS. First, the terminal signals to the circuit switch that it requires a circuit connection to a particular number - this switch is an MSC (mobile switching centre). The MSC has previously downloaded data from the HLR when Mary signed on, into a local database called the visitor location register (VLR) and so knows if she is permitted to call this number, e.g. she may be barred from international calls. If the call is possible, the switch sets up the resources needed in both the core and radio access networks. This involves checking whether circuits are available at the MSC and also whether the radio access network has the resources to support the call. Assuming that the call is allowed and resources are available, a constant bit rate connection is set up from the terminal, over the air interface, and across the radio access network to the MSC - for mobile voice, this will typically be 10 kbit/s or so. Assuming that Mary's dad is located on the public fixed network, the MSC trans codes the speech to a full 64 kbit/s speech circuit (the normal connection for fixed network voice) and transports this to a gateway switch (the gateway MSC - GMSC) to be switched into the public fixed telephone network..

When the call ends, both the MSC and GMSC are involved in producing Call Detail Records (CDR), with such information as: called and calling party identity, resources used, time stamps, and element identity. The CDRs are forwarded to a billing server where the appropriate entry is made on Mary's billing record.

Mary leaves her terminal powered on - so that it moves from being Mobility Management (MM)-connected to being MM-idle (when it was turned off completely, it was MM detached). Mary then boards a bus for the Polytechnic and passes the radio coverage of a number of UMTS base stations. In order to avoid excessive location update messages from the terminal, the system groups large numbers of cells into a location area. The location area identifier is broadcast by the cells in the information they broadcast to all terminals. If Mary's terminal crosses into a new location area, a location update message is sent by the terminal to the MSC and also stored in the HLR.

When Tom tries to call Mary - he is ringing from another mobile network his connection control messages are received by the T-Nova GMSC. The GMSC performs a look-up in the HLR, using the dialed number (i.e. Mary's telephone number) as a key - this gives her current serving MSC and location area, and the call set-up request is forwarded to the serving MSC. Mary's terminal is then paged within the location area - in other words, all the cells in that area request Mary's terminal to identify the cell that it is currently in. The terminal can remain in the MM-idle state, listening to the broadcast messages and doing occasional location area updates without expending very much energy.

Mary and Tom begin a conversation, but as Mary is still on the bus, the network needs to hand over the connection from one base station to another as she travels along. In CDMA systems, however, terminals are often connected to several cells at once, especially

during handover receiving multiple copies of the same bits of information and combining them to produce a much lower error rate than would be the case for a single radio connection. When the handover is achieved by having simultaneous connections to more than one base station it is called soft-handover, and in UMTS, the base stations connected to the mobile are known as the active set.

Mary attends her first lecture of the day on relativity and is slightly confused by the concept of time dilation; she decides to browse the Internet for some extra information. Before starting a browsing session, her terminal is in the PMM (Packet Mobility Management) idle state - in order to send or receive packets, the terminal must create what is called a PDP (packet data protocol) context. A PDP context basically signals to the SGSN and GGSN (Serving GPRS Support Node and Gateway GPRS Support node) which are the packet domain equivalent of the MSC and GMC switches - to set up the context for a packet transfer session. What this means is that Mary's terminal acquires an IP address, the GSNs are aware of the Quality of service requested for the packet session and that they have set up some parts of the packet transfer path across the core network in advance. Possible QoS classes for packet transfer, with typical application that might use them, are: conversational (e.g. voice), streaming (e.g. streamed video), interactive (e.g. web browsing) and background (file transfer). (All circuit-switched connections are conversational.). Once Mary has set up a PDP context, the Session Management (SM) state of her terminal moves from inactive to active.

When Mary actually begins browsing, her terminal sends a request for resources to send the IP packet(s) and, if the air interface, radio access, and core networks have sufficient resources to transfer the packet within the QoS constraints of the interactive class, the terminal is signaled to transmit the packets. Mary is able to find some useful material and eventually stops browsing and deactivates her PDP context when she closes the browser application.

During the afternoon lecture, Mary has her 3G terminal set to divert incoming voice calls to her mail box. Tom tries to ring her and is frustrated by the voice mail - having some really important news about a party that evening. He sends her an e-mail of high priority. When this message is received by the T-Nova gateway, it is able to look in the HLR and determines that Mary is attached to the network but has no PDP context active - it also only knows her location for packet services within the accuracy of a Routing Area (RA). This is completely analogous to the circuit-switched case, and a paging message is broadcast, requesting Mary's terminal to set up a PDP context so that the urgent e-mail can be transferred. Mary is, of course, able to filter incoming e-mails to prevent junk mail causing her terminal to be notified after all, she is paying for the transfer of packets from the gateway.

This scenario has briefly looked at the elements within the UMTS R3 network and how they provide the basic functions of: security, connection management, QoS, mobility management and transport of bits for both the circuit and packet-switched domains. The next section goes into greater detail and expands on some of these points (especially those relating to the packet domain, since this will be contrasted with IP procedures in the next few chapters).

So far, little has been said about the role of the Radio Access Network and the air interface. The Radio Access Network (RAN) stretches from the base station, through a node called the Radio Network Controller, to the SGSN/MSC. The RAN is responsible for mobility management nearly all terminal mobility is hidden from the core network being managed by the RAN. The RAN is also responsible for allocating the resources across the air interface and within the RAN to support the requested QoS.

3.1 UMTS Architecture - Introducing the Major Network Elements and their Relationships

UMTS is divided into three major parts: the air interface, the UMTS Terrestrial Radio Access Network (UTRAN), and the core network. The first release of the UMTS network (Figure 3.2) - R3, the Release previously known as R99 - consists of an enhanced GSM phase 2 core network (CN) and a wholly new radio access network (called the UMTS Terrestrial Radio Access Network or UTRAN).




Figure 3.2: UMTS R3 (Release 99) Architecture.

For readers familiar with the GSM, the MSC, G-MSC, HLR, and VLR (see further reading for more information on GSM) are simply the normal GSM components but with added 3G functionality. The UMTS RNC (Radio Network Controller) can be considered to be roughly the equivalent of the Base Station Controller (BSC) in GSM and the Node Bs equate approximately to the GSM base stations (BTSs - Base Transceiver Stations).

The RNCs and base stations are collectively known as the UTRAN (UMTS Terrestrial Radio Access Network). From the UTRAN to the Core, the network is divided into packet and circuit-switched parts, the Interface between the radio access and core network (Iu) being really two interfaces: Iu (PS - Packet switched) and Iu (CS - circuit-switched). Packet traffic is concentrated in a new switching element - the SGSN (Serving GPRS Support Node). The boundary of the UMTS core network for packets is the GGSN (Gateway GPRS Support Node), which is very much like a normal IP gateway and connects to corporate Intranets or the Internet.

- 3G Base-station (Node B) - The base station is mainly responsible for the conversion and transmission/reception of data on the air interface (Uu) (Figure 3.2) to the mobile. It performs error correction, rate adaptation, modulation, and

spreading on the air interface. Each Node B may have a number of radio transmitters and cover a number of cells. (The Node B can achieve soft handover between its own transmitters (this is called softer handover), the Node B also sends measurement reports to the RNC.

- RNC - The RNC is an ATM switch that can multiplex / de-multiplex user packet and circuit data together. Unlike in GSM, RNCs are connected together (through the I_{ur} interface) and so can handle all radio resourcing issues autonomously. Each RNC controls a number of Node Bs - the whole lot being known as an RNS – Radio Network System. The RNC controls congestion and soft handover (involving different Node Bs) as well as being responsible for operation and maintenance (monitoring, performance data, alarms, and so forth) within the RNS.
- SGSN - The SGSN is responsible for session management, producing charging information, and lawful interception. It also routes packets to the correct RNC. Functions such as attach/detach, setting up of sessions and establishing QoS paths for them are handled by the SGSN.
- GGSN - A GGSN is rather like an IP gateway and border router - it contains a firewall, has methods of allocating IP addresses, and can forward requests for service to corporate Intranets (as in dial-up Internet/Intranet connections today). GGSNs also produce charging records.
- MSC - The Mobile Switching Centre/Visitor Location Register handles connection orientated circuit switching responsibilities including connection management (setting up the circuits) and mobility management tasks (e.g. location registration and paging). It is also responsible for some security functions and Call Detail Record (CDR) generation for billing purposes.
- GMSC - The Gateway MSC deals with incoming and outgoing connections to external networks (such as the public fixed telephony network) for circuit-switched traffic. For incoming calls, it looks up the serving MSC by querying the HLR and sets up the connection the MSC.
- HLR - The home location register, familiar from GSM, is just a large database with information about users, their services (e.g. whether they are pre- or post-pay, whether they have roaming activated, and the QoS classes to which they have subscribed). Clearly, new fields have been added for UMTS - especially relating to data services.

Let us just sketch out the scale of a possible network, taking the UK as an example, to gain a better feel of what it looks like on the ground. First, the Node Bs are the transmitters and will be located in many of the places that GSM transmitters are currently located (site sharing on churches and so forth) - there will also be new sites needed. Many thousands of base stations will be needed to cover 50% of the UK (for example). A short link (maybe microwave) of a mile or so will link the node Bs into something like a local exchange where leased lines connect them to RNCs in regional centers there will be only tens of RNCs. The RNCs are then connected to an SDH ring that is also connected to SGSNs and GGSNs. There will be very few SGSNs, and they will probably be co-located with GGSNs in one or more major centers (combined SGSNs and GGSNs will be available). It is also possible to reuse GSM MSCs and GSNs by upgrading them for 3G.

However, many operators will not want to disturb existing systems and will install new 3G MSCs and SGNs - although these will be collocated with their 2G equivalents.

3.2 UMTS Security

Security in a mobile network covers a wide range of possible issues affecting the supply of and payment for services. Typical security threats and issues might be:

- Authentication - Is the person obtaining service the person who he/she claims to be?
- Authorization - are they authorized to use this service?
- Confidentiality of data - Is anyone eavesdropping on the user's data/conversations?
- Confidentiality of location - Can anybody discover the user's location without authorization?
- Denial of service - Can anybody deny the user service (e.g. sending false update messages about the user's terminal location) to prevent them obtaining some service? An example of this might be when a user is bidding in an auction, and other bidders wish to prevent that user from continuing to bid against them.
- Impersonation - Can users take other users' mobile identities and gain free service, or access to other users' information? Can sophisticated criminals set up false base stations that collect information about users or their data?

In UMTS, there are four main ways in which threats and issues like these are addressed:

- Mutual authentication between the user and the network.
- Signaling integrity protection within the RAN.
- Encryption of user data in the RAN and over the air interface.
- Use of temporary identifiers.

Mutual authentication - of the user to the network and of the network to the user is based around the USIM (UMTS Subscriber Identity Module). This is a smart card (i.e. one with memory and a processor in it), and each USIM is identified by a (different) 15 digit number the International Mobile Subscriber Identity (IMSI) - Note that the IMSI is separate from the phone number (07702 XXXXXX, say), which is known as the Mobile ISDN number and can be changed (e.g. in the recent UK mobile renumbering). When a user switches on, a signaling message is sent to the HLR (their home HLR if they are roaming on a foreign network identified by their IMSI) containing their IMSI and the 'address' of MSC that they are registering with. The HLR (actually in a subpart of the HLR called the authentication centre, AuC) generates a random number (RAND) and computes the result (XRES) of applying a one-way mathematical procedure, which involves a 128-bit secret key (known only to the SIM and the HLR) to the number. The one-way function is very difficult to invert knowledge of the random number and the results of the function do not allow the key to be easily found. The HLR sends this result and random number to the visited MSC, which challenges the USIM with the random number and compares the result with that supplied by the HLR. If they match, the USIM is authenticated. The MSC can download a whole range of keys to store for future use (in the VLR), which is why when a user first turns on their mobile abroad, it seems to take a long time to register but, subsequently, is much quicker to attach. Note that at no time

does the secret key leave the SIM or HLR, there are no confirmed cases of hackers gaining access to these keys in GSM.

A second feature of UMTS is that it allows the user to authenticate the network to guard against the possibility of 'false' base stations (i.e. like bogus bank machines that villains use to collect data to make illegal cards). When the home network HLR receives the authentication request from the serving network MSC, it actually uses the secret key to generate three more numbers - known as AUTN, CK, and IK. The set (XRES, AUTN, CK, and IK) is known as the authentication vectors (Figure 3.3).



Figure 3.3: UMTS authentication.

Both HLR and USIM also keep a sequence number (SQN) of messages exchanged that is not revealed to the network. The MSC sends RAND and AUTN to the USIM that is then able to calculate the RES, SQN, CK, and IK. The USIM sends RES to the network for comparison with XRES - to authenticate itself - but also checks the computed value of the sequence number with its own version to authenticate the network to itself.

Another feature introduced is an integrity key (IK) - distributed to the mobile and a network by the HLR, as described above, so that they can mutually authenticate signaling messages. This takes care of the sort of situation where false information might be sent to the network or to the mobile. This would cover the auction example where a rival bidder sends a false signal that a user may want to detach or have moved to a new base station toward the end of a bidding session.

In addition to the challenge/response, the HLR generates a cipher key (CK) and distributes this to the MSC and USIM. The cipher key is used to encrypt the user data over the air from the terminal to the RNC and is passed to the RNC by the MSC when a connection or session is set up. (In GSM, this key is 54 bits - 54 bits is not that large, and, security-aware readers should note, cracking a 54-bit code is about a one-second job on a custom chip these days.)

UMTS allows the terminal to encrypt its IMSI at first connection to the network by using a group key - it sends the MSC/SGSN the coded IMSI and the group name that is then used by the HLR to apply the appropriate group key. The IMSI is actually only sent over the air at registration or when the network gets lost, and so this new feature should prevent the capture of UMTS identities. After first registration, the terminal is identified by a Temporary Mobile Subscriber Identifier (TMSI) for the circuit-switched domain and a Packet Temporary Mobile Subscriber Identifier (P-TMSI). These temporary identifiers and the encryption of the IMSI at first attach - should prevent IMSI being captured for malicious use and impersonation of users.

One, final, level of security is performed on the mobile equipment itself; as opposed to the mobile subscriber (for example, putting one's SIM in someone else's phone does not always work).

Each terminal is identified by a unique International Mobile Equipment Identity (IMEI) number, and a list of IMEIs in the network is stored in the Equipment Identity Register (EIR).

An IMEI query to the EIR is sent at each registration and returns one of the following:

- White-listed - The terminal is allowed to connect to the network.
- Grey-listed - The terminal is under observation from the network.
- Black-listed - The terminal either has been reported stolen or is not type-approved (wrong type of terminal). Connection to be refused.

3.3 UMTS Communication Management

3.3.1 Connection Management

For the circuit-switched domain, the connection-management function is carried out in the MSC and GMC. Connection management is responsible for number analysis (whether the user is allowed to make an international call), routing (setting up a circuit to the appropriate GMSC for the call) and charging (generation of Call Detail Records). The MSC is also responsible for the transcoding of low-bit-rate mobile voice (10 kbit/s or so in UMTS, the voice data rate is variable) into 64 kbit/s streams that are standard in the fixed telephony world.

The GMSC is responsible for the actual connection to other circuit-based networks and also for any translation of signaling messages that is required.

3.3.2 Session Management

In the packet domain, the user needs to set up a PDP context (Packet Data Protocol Context) in order to send or receive any packets. The PDP context describes the connection to the external packet data network (e.g. the Internet): Is it IP? What is the network called (e.g. BT Corporate network)? What quality does the user want for this connection (delay, loss)? How much bandwidth does the user want (QoS Profile)?

The steps involved in setting up a PDP context are as follows ([Figure 3.4](#)):

- The terminal requests PDP context activation.
- The SGSN checks the request against subscription information received from the HLR (during the attachment). If the requested QoS is not included in the subscription, it may be rejected/re-negotiated.
- The Access Point Name (name of external network) is sent, by the SGSN, to a DNS server (IP Domain Name Server - normal Internet-style name to IP address look up to find the IP address of the GGSN that is connected to the required network).
- The SGSN tries to set up the radio access bearers this can result in renegotiation of QoS.
- The SGSN sends a PDP create context message to the GGSN, and this may be accepted or declined (e.g. if the GGSN is overloaded).

- An IP tunnel is set up between the SGSN and the relevant GGSN - with a tunnel ID (this will be explained in the next section).
- An PDP address is assigned to the mobile.
- The PDP context is stored in the: mobile, SGSN, GGSN, and HLR.




Figure 3.4: PDP context set-up.

In practice, the PDP address will be an IP address (although UMTS can carry X25 and PPP - point-to-point protocol packets as well), and this can be either static or dynamically assigned. In static addressing, the mobile always has the same IP address - perhaps because it is connecting to a corporate network whose security requires an address from the corporate range.

In dynamic allocation, the address can come from a pool held by the GGSN and allocated by DHCP (Dynamic Host Configuration Protocol - again, normal Internet-style IP address allocation) or from a remote corporate or ISP network. The GGSN includes a RADIUS client that can forward password and authentication messages to external servers (as happens in dial-up internet access today). This would typically be the case where users are connecting to their ISPs. So, for example, when Mary begins browsing, she sets up a PDP to Free serve and is greeted by the request for her name and password. These are relayed from the GGSN to the AAA server (Authentication, Access and Accounting) run by Free-serve and, when authenticated, our user's terminal is allocated an IP address belonging to the Free- serve IP address allocation.

UMTS also contains the concept of a secondary PDP context (also called a multiple PDP context Figure 2.8). In GPRS, in order to run two different applications, with different QoS requirements such as video streaming and World Wide Web browsing two different PDP contexts and, consequently, two different PDP (i.e. IP) addresses are needed. In UMTS R99, the secondary PDP context concept allows multiple applications flows to use the same PDP type, address, and Access Point Name (i.e. external network) but with different QoS profiles. The flows are differentiated by an NSAPI (Network layer Service Access Point Identifier, a number from 0 to 15). We will look at the mapping of the various identifiers and addresses later in the mobility management section.




Figure 3.5: Multiple PDP .

A traffic flow template (TFT) is used to direct packets addressed to the same PDP address to different secondary PDP contexts. For example, if a user is browsing and wants to watch a movie clip a long one so they want to stream it rather than download it the browser might activate a secondary PDP context suitable for video streaming. When the video and HTTP packets arrive at the GGSN, they all have the same destination IP address (PDP address). The packet flow template allows other aspects (source address, port number, flow label...) to be used to assign them to the correct context and, hence, QoS. In this case, the source address (or source address and source port number) might be used to differentiate between the flows.

A PDP context will only remain active for a certain length of time after the last packet transmission. In other words, a user might set up a PDP context to browse some web pages and then stop using the terminal. Clearly, they would be tying up network resources (e.g. IP addresses) and almost certainly would not be paying for them (if they pay per packet or by subscription). The network, therefore, deactivates the PDP after a suitable time. It might seem from this that UMTS packet users are confined to user-initiated sessions (the equivalent of outgoing calls only) - but there also exists a mechanism to request users to set up a PDP context. This might be when users have a fixed IP address - so that the GGSN can accept an incoming instant message (for example) and use the IP address as a key in the HLR and obtain the address of the SGSN with which the mobile is associated. When the mobile attached to an SGSN the address of that SGSN was recorded in the HLR as were subsequent movements of the mobile into regions (routing areas) controlled by other SGSNs. The SGSN can send a PDP set-up request to the mobile. Of course, the GGSN has to be careful not to request a PDP context every time a piece of junk e-mail is received. The facility will be more useful when Session Initiation Protocol is used widely for peer-to-peer session initiation.

3.4 UMTS QoS

We saw earlier that when users set up PDP contexts, they included a QoS profile. This section looks at how QoS is described within a UMTS network. UMTS contains the concept of layered QoS - so that a particular bearer service uses the services of the layer below (Figure 3.6). What does this mean? A 'bearer' is a term for a QoS guaranteed circuit or QoS treatment of packets. A concrete example would be that packets leaving the UTRAN - on the I_u (PS) interface - are carried on ATM virtual circuits (that give

guaranteed QoS). Thus, the CN bearer might be an ATM network with virtual circuits offering different QoS characteristics




Figure 3.6: UMTS QoS architecture.

Neither of the local and external bearers is part of UMTS - but they obviously have an impact on the end-to-end QoS. The local bearer might be a Bluetooth link from a 3G mobile phone to a laptop say. In a similar way, the external bearer might, for example, be a DiffServ network operated by an ISP.

At the UMTS bearer level, where PDP contexts are created, all UMTS packet services are deemed to fall into one of four classes (Table 3.2) - basically classified by their real-time needs, i.e. the delay they will tolerate.

Traffic class	Conversational class	Streaming class	Interactive class	Background
Delay	Conversational RT $\ll 1$ s	Streaming RT < 10 s	Interactive best effort Approx. 1 s	Background best effort > 10 s
Example: Error tolerant	Conversational voice and video	Streaming audio and video	Voice messaging	Fax
Example: Error	Telnet, interactive	FTP, still image,	E-commerce,	E-mail arrival

Table 3.2: UMTS traffic classes

Traffic class	Conversational class	Streaming class	Interactive class	Background
intolerant	games	paging	World Wide Web browsing	notification
Fundamental characteristics	Preserve time relation (variation) between information entities of the stream	Preserve time relation (variation) between information entities of the stream	Request response pattern	Destination is not expecting the data within a certain time
	Conversational pattern (stringent and low delay)		Preserve payload content	Preserve payload content

Conversational and streaming classes are intended for time-sensitive flows conversational for delay-sensitive traffic such as VoIP (voice over IP). In the case of streaming traffic such as watching a video broadcast, say - much larger buffering is possible, and so delays can be relaxed and greater error protection provided by error correction techniques that repeat lost packet fragments but add to delays. Interactive and background classes are for bursty, Internet-style, traffic.

When requesting QoS, users invoke a QoS profile that uses the traffic class and seven other parameters to define the requested QoS:

- Maximum bit rate - The maximum bit rate defines the absolute maximum that the network will provide - packets in excess of this rate are liable to being dropped this is equivalent to the conventional peak rate description and is only supported when resources are available.
- Delivery order - The delivery order specifies if in sequence delivery of SDUs is required (for SDU - Service Data Unit read IP packet).
- Transfer delay.
- Guaranteed bit rate - Only the guaranteed rate is always available at all times, and this only applies to the conversational and streaming classes.
- SDU (Service Data Unit) size information - The maximum SDU size.
- Reliability - Whether erroneous SDU should be delivered.
- Traffic handling priority - Traffic handling priority is only used within the interactive class to provide multiple QoS sublevels.
- Allocation/retention policy - Related to the priority of the traffic (this is explained in detail in the UTRAN section later).

There are only certain values allowed for each parameter - more details can be found in the references at the end of the chapter. In practice, however, operators are actually likely to restrict the options for QoS to few basic categories and not try and negotiate all the possible parameters allowed by UMTS.

3.5 UMTS Mobility Management

Most of the mobility management in a UMTS system takes place within the RAN; this was actually one of the design goals of the RAN to hide as much as possible the consequences of users being mobile from the core network. Nearly all handovers fall in this category and have been duly relegated to the [next chapter](#) about the UTRAN.

At the core network level, there are three mobility management states that the terminal can exist in - detached (i.e. switched off), connected, and idle - the last two states having a different meaning in the circuit and packet domains. In the circuit-switched domain, the terminal is always associated with an MSC, and the serving MSC's identity is recorded in the HLR. As described in the example of Mary when a terminal has been idle for circuit-switched traffic for a given time, the network stops tracking it at the cell level, and the terminal simply listens to the broadcast channel of the cells. As it roams about, the terminal is in the circuit-switched mobility management idle mode (MM-idle). Only when it enters a new location area consisting of a large number of cells - does it inform the network of a change of location. When the user wishes to make a call, it performs a procedure called a location update, which provides the network with its position at the cell level of detail. Similarly, if an incoming call is received for the terminal, the MSC broadcasts a paging request for that terminal that immediately responds with a location update, bringing it into the MM-connected state.

Likewise, for the packet mobility management (PMM), when the terminal has not sent or Received any packets for a long time, it ceases to have a PDP context set-up and moves to the PMM-idle mode. When a new PDP context is set up - either as a result of the user wanting to send data or as a PDP context set-up request message - the terminal moves to the PMM-connected state. When a terminal is in the PMM-idle state, it simply listens to broadcast messages and updates the network whenever it passes into a new routing area. (Routing areas are actually subsets of location areas but still comprise many cells).

3.6 UMTS Core Network Transport

This section looks at how data are transported across the core network and how QoS can be achieved. Figure 3.6 shows the user plane protocols for the core and access networks for packet switched traffic.




Figure 3.6: UMTS user plane protocols.

From the terminal to the RNC IP packets are carried in PDCP packets. PDCP is Packet Data Convergence Protocol and provides either an acknowledged/ unacknowledged or transparent transfer service. This choice is related to the (backward) error correction that the underlying RLC (Radio Link Control) layer applies - more details of the functions of RLC can be found in the UTRAN section below. Transparent means that no error correction is applied at Layer 2. The unacknowledged mode detects duplicate and erroneous packets but simply discards them, whereas in acknowledged mode, the RLC operates and resends missing frames (at Layer 2, packets are usually called frames, e.g. Ethernet frames). The choice of mode is based on the required QoS, resending lost or error frames causes delay, and so the acknowledged mode is only used for applications that are delay sensitive. PDCP also performs a compression/decompression function such as compressing TCP/IP headers.

From the RNC to the SGSN IP packets are tunneled using a tunneling protocol called GTP-GPRS tunneling protocol. Another GTP tunnel then runs from the SGSN to the GGSN, allowing a hierarchical mobility (SGSN changes will not happen often) and allowing lawful interception (phone tapping) at the SGSN.

A tunneling protocol consists of a piece of software that takes packets and wraps them within new packets such that the entire original packet including the header becomes the new payload: the original header is not used for routing/switching and is not read whilst encapsulated. A very good analogy is that if a person sends a friend a letter to their home address, their mum puts it in a new envelope, addressed to their college address, and pops it back in the post. GTP in UMTS is more analogous to allowing several languages to be used to address the inner envelopes. Only the main post offices understand Chinese, so letters must be enclosed within a new envelope addressed in English to pass through the UK postal system. Using GPRS tunneling protocol, UMTS can carry a number of different packets (such as IPv4, IPv6, PPP, and X25) over a common infrastructure. GTP packets are formed by adding a header to the underlying PDP packet - the format of this header is shown in [Figure 3.7](#). After forming a GTP packet, it is sent using UDP over IP using the IP address of the tunnel end point, e.g. the GGSN for traffic sent from the SGGN to a external network. The most important header field is the tunnel id that identifies the GTP packets as belonging to a particular PDP context of a particular user (and therefore can be given the appropriate QoS). The Tunnel id is formed from combination of the IMSI and NSAPI - the IMSI uniquely identifying a terminal and the NSAPI being a number from 0 to 15 that identifies the PDP context or the secondary PDP context within a primary PDP context ([Figure 3.8](#)).




Figure 3.8: GTP-tunnelling.

Bits							
Octets	8	7	6	5	4	3	2 1
1	Version			PT	(*)	E	S PN

Bits							
Octets	8	7	6	5	4	3	2 1
2	Message Type						
3	Length (1 st Octet)						
4	Length (2 nd Octet)						
5	Tunnel Endpoint Identifier (1 st Octet)						
6	Tunnel Endpoint Identifier (2 nd Octet)						
7	Tunnel Endpoint Identifier (3 rd Octet)						
8	Tunnel Endpoint Identifier (4 th Octet)						
9	Sequence Number (1 st Octet)						
10	Sequence Number (2 nd Octet)						
11	N-PDU Number						
12	Next Extension Header Type						

Figure 3.9: GTP header format (and formation of packet).

In the UMTS core network, IP Layer 3 routing is, typically, supported by ATM switching networks. It is the operator's choice whether to implement QoS at the IP or ATM level, but if the IP layer is used, the IETF differentiated services scheme is specified by 3GPP as the QoS mechanism. In all cases, interoperability between operators is based on the use of Service Level Agreements that are an integral part of the definition of DiffServ. DiffServ features heavily in the IP QoS chapter, and so readers might wish to read that chapter before looking at the UMTS to DiffServ mapping example below.

If DiffServ is being used to provide QoS in the core network, a mapping is needed at the RNC between UMTS bearer QoS parameters and DiffServ code points, and a similar

mapping is needed at the GGSN for incoming packets. The SGSN originally downloads the user subscription data from the HLR and passes the allocation/retention priority, first to the RNC with the Radio Access Bearer request and then to the GGSN with the PDP context activation message. The RNC and GGSN then use the allocation/retention priority and the UMTS class to map to DiffServ classes. In DiffServ, there is no delay bound, and such a network would rely on proper provisioning to deliver sufficiently low delays for conversational services.

UMTS can support both IPv4 and IPv6 operations and is seen as a key driver for IPv6 technologies. UMTS decouples the terminal packet data protocol from the network transport, through the use of tunneling. As a consequence, it can transport IPv4 or v6 packets without modification. The underlying UMTS core network can also be v4 or v6, and this has no interaction with the user data being tunneled over it.

3.7 Signaling in the UMTS Core Network

The signaling between the mobile, SGSN and GGSN to the HLR, authentication centre, EIR, and also the SMS message centre all consist of SS7 signaling (Signaling System Number 7) messages. SS7 is, in some ways, like an IP network (but it is not IP at all and developed totally independently). It is packet-based and has reliable transport protocols and its own addressing scheme. SS7 was originally used on the PSTN when it became digital and carries all the signaling messages between exchanges need to set up a call (address complete, ringing, and connecting being example messages). The SS7 variant used in the PSTN is called ISUP, and this has been extended for use in mobile networks with an extended message set called MAP (Mobile Application Part). The SGSN/HLR and so forth all have SS7 addresses and use MAP to exchange signaling messages. From the SGSN, the G_r interface connects to a (logically separate) SS7 network over a 2 Mbit/s time division multiplex link (the normal circuit-switched connection that would be found in the PSTN, say, i.e. not IP and totally separated from the data path transmission mechanism).

SS7 is not the only signaling protocol used in the UMTS core network. The setting up, modifying, and tearing down of GTP tunnels are performed by a signaling protocol called GTP-C (whereas the transport of user data is performed by GTP-U, as just described). GTP-C runs between the SGSN and GGSN and also carries the messages to set up and delete PDP contexts. GTP-C uses the same header as GTP-U but is a reliable protocol in that the sequence numbers are used to keep track of lost messages, and these are re-sent. An example GTP-C message is ECHO - this can be sent to another GSN that must reply with an ECHO RESPONSE message that includes the time since the last re-boot. Readers needing more details of GTP-C messages are referred to the TS 29.060 where the nitty gritty detail awaits. There is no SS7 signaling link from the SGSN to the GGSN. Note that GTP-C does not run over the I_u interface between the SGSN and the RNC - since RNCs have no part in PDP context activation, etc. The GTP tunnels from RNCs to SGSNs are set up by part of another protocol RANAP; this is covered in the [next chapter](#).

Chapter 4

UMTS Radio Access Network (UTRAN)

This section looks more closely at the interfaces of the UTRAN, the signaling and transport protocols used to convey bits from the mobile to the RNC, and the underlying ATM switching cloud. The point of this, against the backdrop of a report arguing the merits of IP for 3G networks, is to illustrate several points before embarking on our tour of IP architectures, QoS, and so forth. The first is to explain the switching and timing requirements for soft handover in greater detail, so that the very considerable difficulties of introducing an IP RAN are not overlooked. The second is to demonstrate just how complicated the RAN is, with functions like Radio Resource Management requiring intensive, real-time, processing in a number of distributed elements. RNCs are very large and very expensive. It is not really good enough for IP plaudits to say that IP is simpler and cheaper if, every time a user tries to make a voice call, the quality deteriorates half way through. Finally, the Layer 2/Layer 3 interface is very much integrated and tailored for W-CDMA (with its power control and particular radio characteristics). IP pundits would like a more generic, but less integrated and therefore less efficient, interface that could, say, connect to Wireless LANs as well. For all these reasons, we will wade through the air interface (U_u) and the UTRAN to core (I_u) interface and, finally, tackle the ATM transport.

It is important to note that the standards concentrate on the interfaces because that is where equipment from different manufacturers needs to inter-operate; for example, an Ericsson terminal needs to talk to any make of base station correctly. Within a base station, however, the operation of soft handover is completely propriety, and the standards do not mention this.

4.1 The W-CDMA Air Interface and the U_u Interface

CDMA stands for 'code division multiple access' - meaning that many users share a single block of spectrum by means of different code sequences that they multiply (spread) their data with to increase the bit rate prior to transmission. For example, if a user has a 38.4 kbit/s data stream and spreads it with a chip rate of 3.84 Mchip/s (the spread code bits are called chips), they would have, clearly, a spreading factor of 100. The clever thing about CDMA is that if the spread signal is multiplied by the same spreading code, again, the original bit stream is recovered. Moreover, if there are other users with different spreading codes, the result of multiplying their transmission with the spreading code is simply noise - provided that the codes are carefully chosen. This process is sometimes likened to an international party, where someone hears someone else from the far side of the room talking their language and locks on to that conversation above a general background noise created by other conversations in different languages. One might ask whether it would not be better just to divide the spectrum up and give each user their own part of the spectrum (frequency division multiple access) or time slot (time division multiple access). Many simulations have shown that CDMA can support more users at a given QoS and user bit rate.

In UMTS, there are two modes of operation, FDD and TDD. As has been already explained, the FDD (Frequency Division Duplex) mode uses two blocks of spectrum for the up and down links and separates users solely on the basis of CDMA codes.

The TDD (the time division duplex) mode uses a single block of spectrum but only transmits on the up link or the down link at one time, hence the time division duplex, and uses a mixture of codes and time slots to separate users. The FDD mode will be concentrated on here, because TDD equipment development is lagging behind FDD, not all operators have obtained TDD spectrum, and even those that have, have no immediate plans to exploit it.

In a CDMA system, neighboring cells use the same frequency but avoid direct interference by means of the use of scrambling codes. From the base station to the terminal, the spreading code comprises two parts: the scrambling code that is different for each cell and the chanelisation code that separates users within the cell. Transmissions from users and base stations in neighboring cells are always seen as noise because the spreading code = (chanelisation code \times scrambling code) is unique for each base station to user transmission within the entire system. This is how CDMA is able to use the same frequency in every cell.

In UMTS, the CDMA air interface makes up the physical layer and part of the MAC layer of the U_t interface between the terminal and the base station. The air interface (U_t interface) protocols are shown in Figure 4.1. The PDCP (Packet Data Convergence Protocol) provides header compression for PDP packets - as already described. The BMC (Broadcast/Multicast Control) layer provides cell broadcast facilities - an example of this would be the short message broadcasts of GSM.

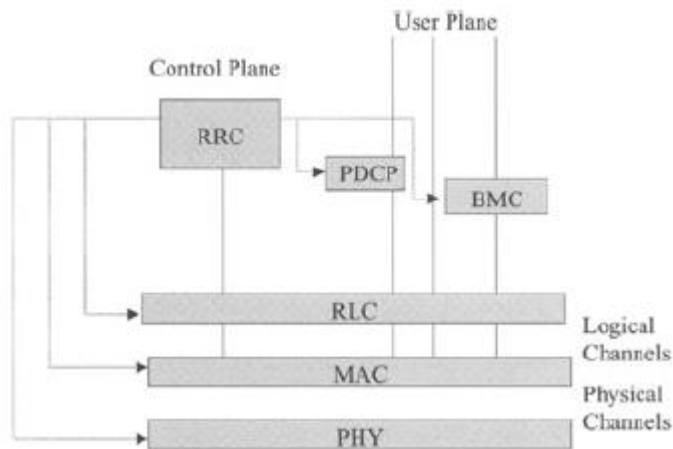


Figure 4.1: Radio Interface protocols - control and user plane.

The RLC (Radio Link Control) layer is responsible for setting up and tearing down RLC connections, each represents a different radio bearer (meaning that there is one radio bearer per PDP context or circuit). The RLC layer segments and reassembles data packets as well as providing backward error correction. A 1500-byte IP packet would be segmented into 27 RLC PDUs with a 2-byte header added to each (the MAC layer would

add another 3 bytes to from MAC PDUs). The level of backward error correction can be one of several modes:

- Transparent - Higher layer packets are not provided with error recovery, and higher layer packets may be lost or duplicated.
- Unacknowledged - This mode detects error packets but simply deletes them. It also avoids duplicating packets.
- Acknowledged - Error-free delivery of packets is guaranteed by an ARQ (automatic repeat request) backward error recovery scheme. Also, duplicate detection ensures that only one copy of each packet is transmitted.

The RLC is also responsible for ciphering and can perform flow control, i.e. the receiving end can request the transmitting end to slow down transmission to prevent, for example, buffer overflow. For data, the RLC terminates at the RNC - so RLC frames are carried to the node B over the radio interface and MAC layers and then on to the RNC on AAL2/ATM switched circuits (see below).

The MAC layer is responsible for mapping logical channels (including data flows) into the transport channels provided by the physical layer. Logical channels in UMTS (FDD mode) include:

- Common control channel (CCCH) - Up link.
- Broadcast control channel (BCCH) - Down link.
- Paging control channel (PCCH) - Down link.
- Dedicated Control Channel (DCCH) - Dedicated (to a single terminal) transport channel (up and down link).

The MAC layer is also responsible for multiplexing/demultiplexing flows from the user on to transport channels (that are similar, but fewer in number to the logical channels, e.g. a BCH (Broadcast Channel carries the contents of the BCCH). The MAC also handles priority handling of flows from one user, i.e. allowing flows with higher priority QoS to have higher priority access to physical channels.

The physical layer is responsible for transmission of data blocks: multiplexing of different transport channels (e.g. the P-CCPCH - Primary Common Control Physical CHannel carries the BCH), forward error correction (error coding) and error detection, spreading (with the CDMA code), and RF modulation.

4.2 UTRAN Mobility Management

4.2.1 Soft Handover

The requirement to support soft handover in UMTS arises from the handover of mobiles between base stations. The boundary between the cells is not clearly delineated and, near the boundary, the ratio of received power from the two base stations fluctuates considerably over even a meter or so (at 2 GHz, the wavelength is 15 cm). If the handover threshold is set such that when the new base station received strength exceeded the old base station strength then the handover would be ping-pong all the time. Each handover 'costs' in terms of signaling messages and network processing time, so, to avoid all the toying and froying, the handover threshold is given hysteresis: once a

handover has occurred, the relative signal strengths must change by 6 dB, say, before another change is made. This is fine for TDMA systems (like GSM), where the interference is felt in distant cells that are re-using that frequency. However, in CDMA systems, having mobiles operating at 6 dB over their minimum power causes a large amount of interference. In CDMA systems, all mobiles interfere with each other, and controlling and minimizing transmit power is the key to increasing capacity. It has been estimated that using the TDMA hysteresis scheme for handover would reduce the efficiency of a UMTS system by 50%. The solution is something called soft handover. In soft handover (Figure 4.2), the mobile receives transmissions from several base stations simultaneously. As the power - and hence the error rate - from each fluctuates the mobile receiver takes the data from each base station and combines them to obtain a reliable answer.




Figure 4.2: CDMA soft handover.

It has been estimated that UMTS mobiles will be in soft handover 50% of the time and will be connected to several base stations simultaneously. For soft handover to work, the frames from different base stations need to arrive at the mobile within about 50 ms or so of each other. Thus, from the network split point, they must be transported to the transmitters with very tight control of delay and jitter. ATM gives this functionality as the transport technology for the UTRAN.

4.2.2 Handover Types

As detailed earlier, in a CDMA system, users are connected to a number of cells, called the active set and cells are added and dropped from the active set on the basis of measurements made by the terminals and reported back to the network. If the cells are served by the same base station (Node B), the mechanism of adding/dropping cells from

the active set is proprietary, i.e. the standards do not specify how it shall be accomplished this is softer handover.

Imagine that a user needs to connect to cells on a different RNC. The original RNC called the serving RNC - connects to the new RNC - called the Drift RNC - via the I_{ur} interface (Figure 4.3). This interface has no counterpart in GPRS or GSM and allows the UTRAN to deal with all handovers independently of the core (this is required in CDMA because of the tight timing constraints for soft handover and the need to add and delete cells from the active sets rapidly). At some point, the UTRAN decides that it should move the SGSN to RNC connection from the Serving to the Drift RNC - a process called SRNS (Serving Radio Network System) relocation. This is essentially a UTRAN function, and the result of the procedure is that the SGSN routes the packets to the new RNC (Figure 4.4).




Figure 4.3: Change of RNC - intra-SGSN SRNS relocation.

If a user moves within the coverage area of a base station (and RNC) served by a different SGSN, this requires the highest level of mobility management - the Inter-SGSN/MSC SRNS relocation (we will concentrate on the packet case). Figure 4.5 shows the situation both before and after SRNS relocation. This is a complex procedure - involving the mobile, 2 RNCs, 2 SGSNs, and a GGSN. The message flows and sequence of events can be seen in 3GPP standard TS 23.060 (downloadable from www.3gpp.org). One noteworthy point about this procedure is that it does not support real-time handover of packets. That is to say, long delay (seconds) can be experienced, and packet duplication is also possible. This is one of the things that will be upgraded in the next release of UMTS. For circuit-switched traffic moving to a new MSC, there is, of course, no gap in service.




Figure 4.4: Change of RNC - inter-SGSN SRNS relocation.

4.2.3 UTRAN Transport

Within the UTRAN, all user data transmission takes place over an ATM-switched network. ATM was originally conceived for fixed networks as a replacement (or evolution) of the PSTN/ISDN to allow packet and circuit data with varying traffic characteristics (constant or variable bit rate) to be multiplexed on to a single connection with guaranteed QoS performance. It was designed to run over optical links with characteristically low error rates (bit error rate of less than 10^{-9}) and so had very light error correction. ATM controls delay and jitter of traffic by carrying all the data inside 53-byte cells (fixed length packets). Because all the cells are the same size, very efficient cell switches could be made that could control the jitter and delay of the cells being switched. Typically, virtual connections would be set up through a mesh of ATM switches (an ATM cloud) by rather slow signaling - hence the term permanent virtual circuits (PVCs) - and voice/IP/video packets or frames or streams would be adapted, using different ATM adaptation layers (AALs), to be segmented and reassembled at the ends of the PVC. The AALs differ in how they segment/reassemble higher layer packets and in the error checking and recovery mechanisms that they provide.

In UMTS AAL2 is used for all circuit and packet data within the RAN, whereas AAL5 is used for signaling within the UTRAN and for transmitting the packet data across the Iu (PS) interface to the SGSN. AAL5 has very little functionality: other than segmenting and reassembling packets into ATM cells, it provides only a basic error check.

The reason as to why AAL2 is used for transport within the UTRAN can be traced back to the particular requirements of CDMA operation and also the desire to support multimedia traffic. 2G's 64 kbit/s circuit switching technology is clearly inefficient for bursty Internet type traffic what is needed is a packet switching technology that multiplexes together many users and uses less underlying (and very expensive) bandwidth in the access network (i.e. the leased lines that typically comprise the UTRAN). The requirement from CDMA, as readers will remember from earlier, was that in order to use a spectrum efficiently, a CDMA system had to exercise very tight power control. This means that it has to support soft-handover with terminals connected to

several base stations simultaneously. More importantly, the combination of signals takes place at layers 1 and 2, and the signals need to arrive within 10 ms or so of each other for correct combination. RNCs need to be able to set up and tear down connections to Node Bs and other RNCs to add and remove cells from the mobiles active set (of cells it is in contact with). This set-up needs to be accomplished rapidly - within 100 ms or so. Finally, the desire to transmit voice, pretty essential for any mobile network, meant that only small packets could be used in the access network. The total delay for a voice call should not exceed about 100 ms. At 12 kbit/s, allowing 20 ms for forming a packet, the packetisation delay - gives a packet size of 240 bits or 30 bytes. An ATM cell carries 48 bytes (compared with an IP packet of typically 1500 bytes - which is one of the reasons why Voice over IP - VoIP - is inefficient). The AMR (adaptive multirate) speech coders used in UMTS networks typically produce a variable bit rate for example; they do not code silence so a constant packetisation delay requires variable length packets. When the coder is only producing 4 bit/s, smaller packets are needed than when it is producing 14 kbit/s - for the same packetisation delay.

From these requirements, along with the need to provide QoS, a new ATM adaptation layer -AAL 2 - and its associated switching and signaling procedures have been developed for the UMTS radio access network. AAL2 allows variable length packets and multiplexes several connections on to a single ATM Virtual Connection. AAL2/ATM is used to carry all the user data - packet and voice - over the UTRAN - the packet data only being converted back to AAL5 at the RNC. The need to multiplex many different, variable rate, traffic sources on to a single ATM VC was the reason for choosing AAL2. Essentially, the key part for UMTS was the development of signaling and switching of AAL2 circuits. A very comprehensive review of AAL2 for UMTS is given in

4.2.4 UTRAN QoS

When a user wishes to make a call or send packets, control signaling for this first passes to the MSC/SGSN. In the packet case, there must be a PDP context active - so that the terminal has an IP address, and the GTP tunnel id is allocated. However, the PDP context does not install any state within the UTRAN, and so each time packet transfer or a connection takes place, radio and UTRAN bearers must be allocated. The SGSN/MSC signals the UTRAN with the required QoS attributes, and these may be granted, renegotiated, or declined. The control protocol (RANAP) used between the SGSN and RNC is described in the [next section](#).

QoS provision within the UTRAN is complicated and can be broken down into air interface and RAN parts. In a well-designed network, the air interface will be the major bottleneck where most congestion and QoS violation will take place. A Radio Resource Management (RRM) function is distributed between the terminal, base station and the RNC and controls QoS over the radio link. RRM consists of algorithms and procedures for the following:

- Admission control.
- Power control.
- Code management.
- Packet scheduling.
- Handover.

In a CDMA system, the common resource consumed as more users are admitted to the system is interference to other users. As users transmit progressively more power, the higher the bit rate or lower the error rate they can achieve, but at the expense of causing higher interference to all other users within the cell and in neighboring cells. The RRM system must ensure that, even after admitting a new connection or packet stream, the overall interference will be low enough to satisfy both the new and existing QoS requirements. Typically, UMTS will run at about 50% of its maximum capacity to ensure stability.

CDMA codes are managed by the RNC - in the down link, channelisation codes must be allocated to each user (so that the spreading code is the cell scrambling code \times channelisation code). The codes for each user must be orthogonal (i.e. have no correlation when multiplied together). In the up link, the scrambling codes are used to separate users and channelisation codes to separate data and control channels from a single user. All these codes are allocated/de-allocated and managed by the RNC.

Packet scheduling takes place each time a user or a base station has packets to transmit (other than the very small amount of data that can be transmitted on the random access channel). A request is made to the RNC and only admitted if resources are available. These requests can be queued - and the way in which they are treated depends on the allocation/retention parameter. This is used by the RNC when deciding how to allocate resources when faced with new QoS resource requests. The idea here is that operators can offer different priority levels for resource allocation - even for users who have requested the same QoS. As an example, a gold user (business class) may pay a large subscription that result in an appropriate allocation and retention policy entry in the HLR against the user's name. When the user wants to transmit some packets, a request arrives at the RNC and is treated with the respect it deserves, i.e. it goes to the head of the request queue, and even if no resources are available, it causes lower-priority users such as bronze (economy class) - to have their resources reduced. In real terms, the gold user is able to make their important VoIP call to the chairman of selectors at Lords, and some impoverished student loses half the bandwidth of their video download.

The allocation/retention parameter is complex and contains information on: priority, preemption capability, pre-emption vulnerability, and whether the request can be queued. The resource scheduling algorithm, located in the RNC, is not specified by 3GPP and is vendor implementation-specific. The final reply from the RNC is either success/failure or failure due to timing out in a request queue.

One proposal for a practical, flexible way of providing QoS to users for Internet services is the concept of services classes. Let us assume that there are three classes - gold, silver, and bronze. Each class offers certain group behavior to users of that class. An example would be:

- Gold users always obtain the requested bandwidth - regardless of interference, congestion, or radio degradation (unless they themselves are already using all the available resources).
- Silver users have an elastic bandwidth, but the service is better than that experienced by bronze users, who obtain the equivalent of a best-effort service.
- Bronze users have only best effort QoS.

Of course, different operators will operate different QoS schemes, and the standards do not specify exactly how QoS is achieved or implemented in UMTS, just how it is signaled across the interfaces.

The radio bearer is also responsible for mapping the appropriate class to the correct error correction mechanism. Conversational services tolerate little delay and need forward error correction (redundancy added to the transmitted frame), as opposed to backward error correction (re-transmission of corrupted/lost frames), which causes delay.

Handover in UMTS is handled by the RRM. As we have seen earlier, this includes softer handover (when handled within a base station) and soft handover, where the terminal has an active set of base stations that it is in contact with. The RRM decides when base stations are added or deleted from the active set, sometimes, this is suggested by the terminal, but the network always has control of this process.

4.2.5 UTRAN Signaling

The signaling across the Iu interface (from SGSN to RNC) is provided by RANAP- Radio Access Network Application Part. RANAP is responsible for:

- Radio Access Bearer set-up, modification, and release.
- Control of the UTRAN security mode.
- Management of RNC relocation procedures.
- Exchanging user information between the RNC and the Core Network.
- Transport of mobility management and communication control information between the core network and the mobile [the so-called Non-Access Stratum (NAS) information - such as PDP context management - that does not concern the UTRAN].
- Set-up of GTP tunnels between the SGSN and the RNC.

From the RNC to the terminal, the Radio Resource Controller (RRC) sets up a signaling connection from the user's equipment (UE) to the RNC. This covers the assignment, re-configuration and release of radio resources. The RRC also handles handover, cell re-selection, paging updates, and notifications.

4.3 cdma2000 Packet Core Network

Cdma2000 is another member of the IMT-2000 quintet of 3G standards and has a North American origin. Originally, there was cdmaOne - which utilized the IS-95A CDMA air interface and an ANSI-41 network (ANSI-41 is another, GSM-like, network of base stations, base station controllers, and specifies the protocol used to signal between them). CdmaOne was launched in Hong Kong in 1995 and is now widely used for voice and low-bit-rate data in the Far East and North America. The first upgrade to cdmaOne was a new air interface - IS- 95B with data rates of up to 64 kbit/s being offered by some operators today (Q2 2002). Cdma2000 comes in two stages - 1X and 3X. cdma2000 1X is designed to be backwards compatible with cdmaOne and offers twice the voice capacity (in the same spectrum) and data rates of up to 144 kbit/s. cdma2000 3X pushes the maximum data rate to 2 Mbit/s and offers a greater efficiency than 1X. Unfortunately, cdma2000 is not compatible with UMTS - the need for backward compatibility means

that two CDMA systems use completely different bit rates, frequency blocks (cdma2000 uses multiple carriers in the down link for cdmaOne compatibility), etc.

The packet core network (PCN) is a network architecture being promoted by the US standards group TIA (Telecommunications Industry Association) for cdma2000 networks: the overall architecture is shown in [Figure 4.5](#).




Figure 4.5: PCN - Packet Core Network architecture.

Much like UMTS, this includes a radio access network (and all the concomitant issues of soft handover) as well as a PDSN (Packet Data Service Node), roughly equivalent to a SGSN. The major difference between the PCN and UMTS is in the way that mobility management is handled. In UMTS, as we have seen, this is handled in the HLR and uses SS7 signaling the PCN, however, is based on Mobile IP (MIP) an Internet mobility concept.

MIP has been developed in the IETF because operating systems and applications, such as FTP, are not tolerant to a change of IP address - many operating systems require a complete restart. Current IP routing protocols would need modifying to allow users to roam with a constant IP address; they would need more processing power, much larger tables containing so-called per-host entries (i.e. one per user), and a protocol to keep track of the user's movements. Such protocols exist but are very much at the research stage. A less elegant solution is to introduce a fixed point - the home agent (HA) - which receives all packets sent to a users home address. When the user is roaming on a foreign network, they are allocated a care-of address by a foreign agent (FA) (a process that sends out advertisements and responds to requests for addresses). The roaming terminal then tells its home agent what its care-of address is, and when packets arrive, addressed to the home address, the HA tunnels them - using IP in IP tunneling - to the FA. The FA de-capsulate the packets and sends them on the mobile user using a Layer 2 address (remember that the FA and mobile are on the same subnet). Packets from the mobile to correspondent host (CH) can be sent directly - since they have a destination address that routers can use directly. This is MIP and the great beauty of it is that it does not require any change to existing routers, routing protocols, or existing IP stacks that are not mobile aware. The best analogy is the postal one already described in tunneling I send my letters

(packets) to your home address and your mother (home agent) puts them in a new envelope (encapsulates them) addressed to your college address (care-of address). When the letters arrive you (acting as a foreign agent) tear open the envelope and find the original letter inside. You keep in touch with your mum by separate letters (signaling) that include notification of your new address (registration) and your mum knows it is you because she recognizes your handwriting (authentication!).

The PDSN acts as a foreign agent - providing care-of addresses to mobiles and de-capsulations IP packets tunneled from the HA. The link from the PDSN to the mobile is made using PPP (point-to-point protocol), PPP is more familiar from dial-up networks, where it provides encapsulation and error protection from computers to the NAS (network access server). The radio interface between the RN and PDSN has two channels one for data and the other for signaling. The signaling comprises standard MIP messages (registration request, registration reply) plus two additions registration update and registration acknowledge. Terminals, therefore, have to run a cdma2000 special IP stack containing the non-standard MIP code. MIP does nothing to provide mobility in the radio network, and, since this must support soft handover, this is based not on IP but on ATM. The additions to the standard MIP framework are:

- The use of an AAA (Authorization, Authentication, and Accounting) server.
- The packet data-related RADIUS attributes.
- Use of the PDSN node as the FA.

The AAA (Authorization, Authentication, and Accounting) server is a standard IP server carrying details about subscribers and is used to authenticate users and check that they are able to use the requested service; in this respect, it is performing a similar role to the HLR UMTS. It also stores and forwards accounting information usage data records generated by the PDSN.

The PDSN both acts as a mobility anchor and operates as a RADIUS client in forwarding authentication details towards the appropriate AAA server. The RADIUS protocol Remote Authorization Dial in User Service has been extended to carry additional attributes specific to cdma2000 typically session status, differentiated service class options and accounting attributes.

A typical MIP registration in cdma2000 would comprise the following:

- A mobile roams to a new network and sends a request for service to the base station; the base station checks for existing links and then forwards the request to the PDSN.
- The mobile negotiates with the PDSN using PPP and sends its id and security data to the PDSN.
- The PDSN uses the id to locate the home AAA and places the security data into a RADIUS request.
- If the home AAA authenticates the user, then the PPP link is finally established.
- The mobile node then solicits for a FA - the PDSN responds with advertisement showing available foreign agents.
- The mobile creates a MIP request and obtains a FA care-of address.
- The mobile forwards its care-of address probably securely to the HA, which creates a tunnel for any incoming packets.

If the mobile moves from the area of one base station controller to another, and they are both connected to the same PDSN, the PDSN is then able to associate the mobile with the old PPP state this is an intra-PDSN handover. If the mobile moves into the area covered by a new PDSN, it has effectively roamed to a different foreign network and must establish a new PPP connection, obtain a new care-of address, and register again with its HA.

4.4 Conclusion

3G was conceived in the late 1980s to meet the 'Martini' vision – communications 'anytime, anyplace, anywhere'. However, the 3G concept has undergone radical changes since then – the satellite component, the B-ISDN, and the fixed mobile convergence ideas have all been removed. The industry - operators and manufacturers - have also chosen an evolutionary route for 3G core networks, leveraging their extensive investment and research in 2G networks. The reasons for this include the runaway success of 2G voice networks - to the point where costs are now rivaling fixed-line operations for voice traffic. GSM and other 2G technologies were the products of a tight standardization process - where complete, integrated, solutions were specified by means of interfaces. This process has also been adopted for 3G standardization.

In sharp contrast, the air interface(s) chosen are revolutionary, and the decision appeared to be, to many neutral observers, political fudge that allowed five different standards to be called 3G. The nature of CDMA, requiring soft handover support, and the need to multiplex variable rate voice and multimedia traffic have also led to the adoption of the transport technology of the 1980s - ATM - for the radio access network.

Whilst 3G has been moving through the standardization process, a second great revolution, after 2G mobile, has been unfolding - the Internet. Nowadays, nearly all data transmissions, and an increasing number of voice calls, are encapsulated within IP packets before leaving the end terminal. IP services have grown rapidly, with e-mail, browsing, and countless business applications being built on IP technology.

There is no doubt that 3G will have to carry IP packets to be successful. 3G is not about voice - 2G networks are highly optimized for voice and can be increased in capacity by using smaller cells or lower rate voice coders. 3G has to be about multimedia, where the traffic is bursty and the bandwidth variable and (potentially) much higher than in today's 2G networks.

Without doubt, traffic on such a multimedia network will be predominantly IP-based. 3G can certainly cope with IP multimedia - at one level, it can be simply viewed as a non-IP access technology like PSTN dial-IP access or GSM circuit-switched data. For example, UMTS can be easily viewed as a Layer 2 network, since user IP packets are always encapsulated, and the headers are not used until the Internet gateway is reached. 3G has features like multiplexing in the RAN and bandwidth on demand in the air interface to support this bursty packet traffic and will offer data rates from 64 kbit/s to something like 200–300 kbit/s.

The real issue with IP traffic over 3G, however, is one of efficiency, cost, and revenue earning potential. [Chapter 1](#) has detailed some of the reasons for why 3G might be inefficient when it comes to carrying IP traffic, including the high cost of providing advanced IP services such as multicast, the inability to evolve to meet the rapidly changing open IP architecture, the economic cost of an inefficient transport of IP packets in the core network, and, finally, the lack of support for IP applications within the mobile network. Web servers, for example, are located outside the mobile network.

[Chapter 8](#), will examine the arguments for a more IP-based 3G network, sketch out the design for an all-IP mobile network, and examine how 3G is evolving to take into account the development of IP, and other, technologies. First, however, it is time to look at research and designs for IP solutions to the functions required to build a mobile network: security, mobility management, and QoS

Chapter 5

Multimedia Service Support and Session Management

5.1 Introduction

Two of the key new features of 3G networks are their ability to support multimedia applications and the Virtual Home Environment. The former implies a network with the ability to support more than just voice communications (and more than just non-real-time, data applications like the World Wide Web and e-mail). The latter is where users of 3G networks store their preferences and data. In its original sense, as described in [Chapter 2](#), the VHE is responsible for tailoring the communications to the physical connection and terminal currently being used. This chapter considers how this type of functionality could be provided in an IP network. It begins with a discussion of the key concept of session management. A multimedia communication, such as a video-telephony call, is referred to as a session. There are a number of different functions that are required to provide and support sessions. This chapter focuses particularly on the session management control plane functions. Following this, we briefly consider how currently sessions and VHE functionality are handled in both 2G/R99 UMTS systems and the Internet. Within the Internet, control plane session management for real-time, multimedia services is an area that is still under development. The two main protocols for this role are reviewed. H.323 is currently in use today, whereas the Session Initiation Protocol (SIP) is a newer IETF standard. SIP is included in the next generation of UMTS standards. Its operation is then examined in some detail. The chapter then goes on to look at some examples of the power of SIP, how it could be put to use in 3G networks, in particular, how it can be used to link between traditional telephony networks and IP networks, and how SIP can enable advanced networking services. Throughout this chapter, SIP is considered in the context of a future, mobile, multimedia Internet. The use of SIP in forthcoming versions of UMTS is rather different to this model - the 3GPP additions to SIP make it almost an entirely new protocol altogether.

As SIP becomes better understood, it will become clear that, in addition to its role in multimedia service support, SIP is highly related to the original VHE concept.

5.2 Session Management

5.2.1 What is a Session?

A session is a series of meaningful communications between two or more end points. Sessions are supported by connections (such as a TCP /IP connection) that provide the physical connectivity, which ensures that bits flow correctly between the end points. The session provides the additional support that enables the receiver(s) to determine whether a particular stream of bits should actually be transformed into an audio-stream, for example.

A session may have many connections associated with it. An example of this is a video conference, where the audio and video parts of the data are sent over separate

connections. Further, a single connection may remain active through the lifetime of several sessions.

5.2.2 Functions of Session Management Protocols

Session-layer (signaling) protocols are used for creating, modifying, monitoring, and terminating sessions with one or more participants. These sessions include multimedia conferences and Internet telephone calls.

To illustrate this, consider a typical procedure that would have been required to establish an Internet Voice Call more than 7 years ago, running between two users at adjacent desks. The two users would first ensure that they would both, be using the same application, agreeing on the nature of the voice coding, sampling rate, data compression, and error coding that would be used. IP addresses would be exchanged, and UDP may have been agreed on as the transports control mechanism, so that the connection could be established. At this point the users would stop talking and actually boot up their computers. Today, this entire process is part of 'Session Initiation' or 'the control plane of session management', and a number of different protocols exist to facilitate this process. This process is studied in depth in this chapter.

Typically, on a first attempt at an IP voice call, speech would be much distorted because other traffic on the local Ethernet would be causing severe, variable, packet delays. Packet delay is very important for any real-time communications and can be heard as the very awkwardness often associated with television interviews carried out over satellite because of the considerable length of time between the interviewer asking a question and the interviewee responding. For good communications, the end-to-end delay needs to be no more than about 150 ms. There are several sources of delay: packetisation delay, transit delay, queuing delay, and buffer delay. Packetisation delay is the time it takes to fill a packet, and 20 ms is considered the usual upper limit. This is why data packets containing voice are often very small. The transit delay is simply the minimum time that it takes the packets to be transmitted physically across the wires and processed by the routers. Within the Internet, this can vary from packet to packet with the route taken. Queuing delays are the variable delays at the routers caused by other traffic sharing the router (or, in our example, the variable delays caused by our packets waiting to get on the Ethernet along with large packets associated with file transfers). The buffer delay is how long the packets wait in the buffer at the receiver to be played out. This is a tradeoff, as longer buffer delays allow more packets to arrive and so reduce the number of lost packets, which also affects speech quality. Much of the work on Quality of Service, discussed in [Chapter 7](#), is concerned with tackling the problem of queuing delays. This requires co-operation between the end terminals and the network.

If packets are played out as soon as they arrive at the terminal, then any variability in the delay (known as the jitter) compounds the problem of speech distortion. To overcome this problem, the Real-Time Protocol, RTP, and the associated Real-Time Control Protocol, RTCP, are typically used within the Internet. These are session layer, end-to-end protocols that do not require any co-operation from the network. They ensure that packets within a session are played out at the correct time. As well as overcoming the problem of jitter, this is particularly useful when a session consists of multiple

connections (audio and video), because these need to be correlated so that the speaker's mouth is seen to open when they start to speak. Although RTP and RTCP are (data plane) session management protocols, they directly affect the quality of the communications; they are discussed further in [Chapter 7](#). Without RTP/RTCP, earliest attempts at Internet telephony only achieved satisfactory performance if the two machines were directly connected, for example with a dedicated Ethernet.

5.2.3 Summary

A session is a multimedia communication, where 'communication' implies some sort of semantic understanding and is distinct from the connection and transferal of bits. Sessions are important concepts in both supporting multimedia applications and in providing the VHE of 3G systems. This chapter will focus on control-plane session management protocols. The key functions required by such a protocol are:

- Locating the parties to be involved in the session.
- Negotiating the characteristics of the session.
- Modifying the session.
- Closing the session.

A session management protocol should automate much of this procedure essentially leaving a background process listening on a fixed port on the terminal to handle such requests and alerting a suitable peer application. Further, such a protocol should be able to support multiparty calls. The application may use information about local resources and their understanding of the network to negotiate the session characteristics. An example of this would be an application that knows it has a wireless network connection and so suggests a low bit-rate voice encoding. Once the session is established, the receiver, using RTCP, will normally identify serious QoS violations. The session control protocol will then allow the terminals to change the session description to match the available resources. Ideally, the session protocols should give the sender sufficient information so that, should it detect a QoS violation, it knows how to adapt its data.

'Session' is a highly generic term and is used in different ways in different communities for example, the term 'connection' used in this report will be called by others 'a session at the transport level'. We have tried to avoid this confusion by defining our terms, but the reader should be forewarned that not all texts use the same definitions.

5.3 Current Status

5.3.1 Session Management

Session management functionality seems so essential, but session management today often goes unnoticed. Essentially, whilst 'session' is a generic term that includes everything from real-time multimedia communications to a simple web download, explicit session management is currently only considered in the context of multimedia and/or real-time communications. The reasons behind this will become clearer in the following sections that look at how sessions are managed in today's networks.

5.3.2 Within 2G Networks

Traditional circuit-switched telephony networks only support one service voice. A voice session is typically known as a phone call. The data rate and encoding schemes are clearly defined, and special inter-working units - media gateways - need to exist to translate data dynamically between the encoding schemes used in different systems (e.g. between the PSTN 63 kbit/s networks and 2G 13 kbit/s networks). Session management and quality of service are tightly integrated within the application and network. Features like session divert (where an incoming phone call can be redirected from the office to the mobile phone) and call (session) waiting are provided using dedicated, specialized platforms known as Intelligent Network (IN) platforms.

This approach works well for a single service. There is no overhead in negotiating a session. The network can easily provide service quality, using Erlang's formula, to dimension resources. However, it becomes very difficult to support multimedia services in this way. One issue, for example, would be the number of types of translation that a media gateway would need to be able to perform. The development of services in the Intelligent Network platform is also complex and time consuming.

In 2.5G, GPRS, there is still no concept of an explicit session, and again both session management and quality of service management are tightly coupled. Users set up a PDP context and connect to their access network provider- an ISP or corporate LAN. They can access services such as web browsing and e-mail, but real-time interactive services will not be supported. Also, multicast services will not work because of the use of GTP.

5.3.3 Within the Internet

Mail and web browsing are the most commonly used Internet applications. Here, web browsing will be considered as an example of current session management. In essence, there is only one type of web download - the user finds the machine and takes the data using TCP to provide reliable data transport. The data come across as plain text, which is then displayed in the browser. It is a 'one size fits all' approach. In fact, DNS is used to find the IP address to enable a connection to be established to the correct web server. MIME types (originally developed for mail, but extended to be applicable to the web) then provide some form of session information, telling the browser what type of data will be received. However, there is no negotiation of this information - the user cannot choose a 'gif' over a 'jpeg' version of a file - the file is already written and stored on disk. Thus, some session management functionality is already available as a very familiar protocol, and the rest of the required session management is incorporated within the basic HTTP web protocol. This approach works well when there is a limited amount of session information that needs to be exchanged.

5.3.4 Session Management for Future Applications

Multimedia and real-time sessions are much more complex. There are many more parameters (such as error coding schemes and data rate) to agree on - at least if the user wants to ensure that the quality of the session is good. There are more parameters partly because it is harder to achieve good quality for real-time communications than for a web session. With web, data should be accurate and fairly timely. With a multimedia session, a user may trade, for example accuracy for delay, or a low-resolution video for a high-resolution audio stream. Also, data are not yet encoded, so there is a chance for the user to choose the best data format for their terminal and network. There may be a whole range of different applications that would be able to inter-work if only this information could be negotiated. Thus, it makes sense to abstract the generic session initiation functionality, and provide a protocol that can be reused by many different applications. Such a protocol would promote connectivity, which was previously argued as key for the growth of the Internet. Further, although DNS enables us to find computers, for real-time communications, we are often more interested in finding a person to talk to. Some applications (particularly Instant Messaging applications, such as ICQ) have provided their own systems for locating users. In this situation, the user can register their permanent identifier (*your.name@chatserver*) at a central server, together with the IP address of your current terminal, and start a process (application) on their machine that listens on a particular port. When somebody wants to contact the user, they can send a message to the server that is then able to tell if the user is on-line and deliver the message, confirming delivery to the sender. However, again, it makes sense to have a generic, reusable system for the function of locating users

5.4 VHE Concept

The original VHE concept has previously ([Chapter 2](#)) been described as:

Where users of UMTS would store their preferences and data. When a user connected, be it by mobile or fixed or satellite terminal, he or she was connected to their VHE which then was able to tailor the service to the connection and terminal being used. Before a user was contacted then the VHE was interrogated - so that the most appropriate terminal could be used and the communication tailored to the terminals and connections of the parties. Thus, there is a close relationship between session management negotiation of a session's characteristics and the VHE concept.

5.4.1 Within 2G/3G Networks

The VHE concept in 3G networks has been reduced to the GSM equivalent CAMEL (Customized Applications for Mobile network Enhanced Logic). CAMEL is a GSM specialized IN platform that allows users to roam on foreign networks and still receive some of the advanced services that the home network operator provides. These are all switched circuit and voice-based, and a good example is short code dialing for voice message retrieval. In the UK, users can dial 901 to obtain messages; in France, this does not work, but CAMEL intercepts the dialed number and queries the home HLR to allow number substitution (just like fixed network IN), giving the French switch the correct number 0033563867387 (say).

CAMEL is about more than just standardized IN services, however. It is designed to support flexible service control and creation, so that operators can quickly deploy advanced value-added services. These services can be accessed by a user, even if they are roaming. CAMEL enables this by providing a standardized interface between the network entity controlling the new services (called the GSM Service Control Function, GSM SCF) and the visited network's switches.

Figure 5.1 shows the generic architecture for CAMEL. Apart from the standard GSM elements (HLR, MSCs, and VLR), a new entity has been introduced: the CAMEL Service Environment (CSE) - that encompasses the gsmSCF. New functionality has also been added to the mobile switches: the gsmSSF (Service Switching Function).

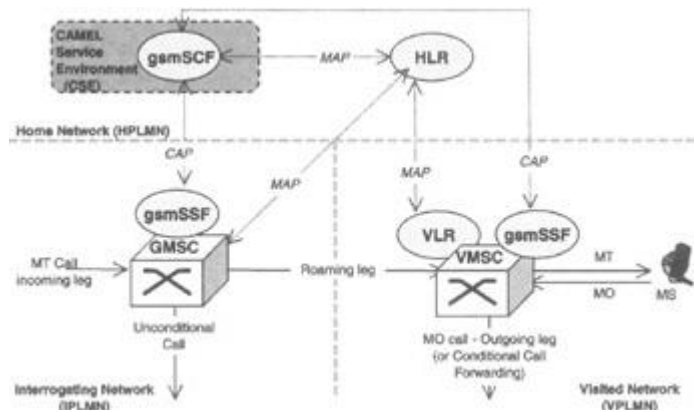


Figure 5.1: Functional architecture for support of CAMEL. GMSC: Gateway Mobile

Switching Centre, VMSC: Visited MSC, VLR: Visited Location Register, HLR: Home Location Register, MAP: Mobile Application Part, MS: Mobile Terminating, MO: Mobile Originating, SSF: Service Switching Function, SCF: Service Control Function, CAP: CAMEL Application Part, CSE: CAMEL Service Environment.

CAMEL is being extended for use in later releases of UMTS - including PS domain and IP telephony capabilities. The interface between the CSCF and the CSE is still being discussed within 3GPP. The IM domain will, then, have options for SIP, CAMEL, and a PARLAY-style interface for service creation. The PARLAY-style interface will be based upon the OSA (open service architecture) being specified by the OSA group within 3GPP. However, CAMEL follows a very different model to that of Internet services. The service provider is still the network provider. The services being managed are still just voice services.

5.4.2 Future VHE

Internet Portals provide the closest service to the VHE that can be seen in the Internet today. The reader may be familiar with them - they are the websites that ISPs encourage customers to have as their home page. Being web-based, they can be accessed from any terminal. Everything can be accessed, from mail to daily newspapers, from these sites. However, neither the first generation of UMTS networks, nor the Internet can provide the VHE functionality as originally described in early UMTS visions. The concept of the VHE will be revisited in the final section of this chapter.

If you feel we are mixing our layers here - it is very easy to do in telephony style networks, where everything is tightly integrated.

5.5 Session Initiation Protocols

Previous sections have highlighted what session initiation protocols are required to do to find a user and enable multimedia communications to be established. Once the session is running, RTP and RTCP (both well-known, stable protocols) are used to manage the session. However, the protocols for session initiation, the ITU H.323 and the IETF Session Initiation Protocol (SIP) - are much less stable, and still under development.

In considering these session initiation protocols, attention is focused on multimedia and real time applications, as these are the applications where generic session management protocols will give the greatest benefit.

5.5.1 H.323

The H.323 protocol suite is a full session control protocol it includes session creation, data transport, and data plane session control functionality (the latter through RTP). This protocol was originally developed in the early 1990s and is standardized by the ITU. It was initially focused on videoconferencing and is currently integrated into a number of applications including CUSeeMe Professional and Microsoft's Net-meeting. However, perhaps as an indication of the complexity of the standard, only recently have these two standard compliant solutions been able to inter-work.

The current standard has a number of weaknesses however, making H323 more suitable for LAN environments than the Internet. One of the most significant issues is the fact that it is a heavyweight protocol. For example, establishing a session using H.323 can take 7 round trip times. The signaling must be transported using (multiple) TCP connections, which is an unnecessary overhead for wireless applications and also complicates the implementation of firewalls. It also includes a large amount of functionality that is available already through other Internet standards - it is less a modular than a stove pipe solution. It requires state to be held through the network, making it less suitable for wide area networks. Finally, user mobility can lead to routing loops. H.323 is still under development to tackle these criticisms. The next version (3) should include fast call set-up and UDP signaling, and should solve the routing loops, but is not yet available as a standard. There is some evidence that H.323 will eventually converge with its new rival, SIP, but convergence is slow. Whilst it is widely used in applications, there is less evidence of it being widely supported by network operators (the operator support is required for large-scale networks and directory services).

5.5.2 SIP

The Session Initiation Protocol (SIP) is a much more recent development. It was originally developed between 1996 and 1999 in the IETF MMUSIC group and at Colombia University. The SIP IETF working group was formed in September 1999, and a draft standard of SIP appeared in July 2000 from the IETF. It is a general, multimedia, session initiation protocol. It is smaller than H.323. It is transport layer independent - although most implementations use UDP transport. It is lightweight; for example, it only requires 1.5 round trip times to establish a session. By using UDP, it simplifies

multicasting, which facilitates applications such as user location at a range of terminals or call centre applications. Unlike H.323, it does not specify anything about resource reservation or security - other protocols deal with these aspects. It is the view of many within the IP community that this limited scope of SIP is precisely the aspect of SIP that makes it so powerful.

SIP is a text-based protocol, similar to HTTP. Such systems tend to be easier to debug and integrate with high-level programming languages.

SIP also allows far more extensive error and status reports than H.323. SIP is almost invariably used to carry session description messages, as defined by the session description protocol SDP but even this is flexible. To allow for fast adaptation, several SDP objects could be agreed upon in session initiation. As well as being a simpler protocol, SIP is regarded as more general. It can operate in end-to-end and proxy server modes, and it supports both distributed control and centralized bridge architectures for multiparty calls.

5.5.3 Session Initiation for 3G

H.323 came first, so developers of SIP could learn from the H.323 experience. This has resulted in SIP being both a simpler and more flexible protocol. The mapping from SIP to H.323 is relatively easy and well defined, whereas the converse is not true. Thus, 3G networks have decided to use SIP rather than H.323, so SIP will now be discussed in more detail. Its memory footprint, and also a rough word count of the relevant standards documents.

5.6 SIP in Detail

5.6.1 Basic Operation of SIP

The Session Initiation Protocol (SIP) is a means of negotiating contact between one or more entities, whether they are individuals or automats. On its outward face, SIP manifests itself as an application - the User Agent. The SIP messages are few and entirely in plain text, requiring very little processing. They are rich and readily extensible. Media negotiation can be included within SIP messaging, utilizing Session Description Protocol (SDP) or MIME types (or anything else) within the body. SIP itself is not a data carrier; other protocols such as UDP do that. SIP is solely the means of negotiating contact and exchanging the necessary parameters to trigger applications. SIP specifies six methods for *initiating* contact, the most common of which is the INVITE method. User Agents are required on each of the participating machines ([Figure 5.2](#)).




Figure 5.2: SIP signaling during call set-up.

In this simple scenario, User Agent A is being used to initiate contact with B. User Agent B's IP address is known in advance, so User Agent A simply opens a socket and sends an INVITE message to the destination. Note that both User Agents are listening on port

5060: this is the default port for SIP. User Agent B receives the invitation, and now has to return a RESPONSE from the many defined by SIP. In this case, the invitation is accepted by returning OK. Other RESPONSES (from about 30) include: BUSY, DECLINE, and QUEUED.

The format of the SIP message is twofold: a header, consisting of SIP fields, and a body. Header fields provide such parameters as the identity of the caller, the identity of the receiver, a unique call id, sequence number, subject, the hop traversed to deliver the message (i.e. VIA), and so forth. The body typically uses SDP to describe the session that is being negotiated. In the above example, User A might specify that they wished to invite B into a media session, including audio (Figure 5.3).

```
sip:mel@uk.net SIP/2.0
Via: SIP/2.0/UDP spain.tel:5060
From: Clive Dellard <sip:clive@sipstreme.net>
To: Mel Bale <sip:mel@uk.net>
Call-ID: 10000001@sipstreme.net
CSeq: 1 INVITE
Subject: Urgent Call
Contact: Clive Dellard <sip:clive@sipstreme.net>
Content-Type: application/sdp
Content-Length: 160

v=0
o=clive 4534593492 3284729843 IN IP4 sipstreme.net
s=Session SDP
e=clive.dellard@bt.com
c=IN IP4 100.101.102.103
t=0 0
m=audio 9160 RTP/AVP 0
a=rtpmap:0 PCMU/8000
```

Figure 5.3: Typical SIP INVITE message.

SDP provides fields to specify the intended applications, codec's, and endpoint addresses. If B can support A's suggestions, B simply copies the SDP body back to A in his OK RESPONSE, entering his own endpoint addresses and port numbers for the medium. Thus, session negotiation and set-up can take a minimum of three SIP messages, i.e. just 1.5 network round trips. However, should B not support one particular codec, but can offer another, they would amend this field in the SDP of their returned OK. If the change is acceptable to A, the ACK follows as normal; otherwise, A CANCELS the session, or re-negotiates, sending another INVITE, with a new SDP, but the same Call ID and a higher sequence number. B recognizes the Call ID and realizes that it is a renegotiation from the earlier sequence number, and the process begins again.

In the same way, in-session re-negotiation is supported, e.g. the existing video session is streaming, and A decided to add voice. The other SIP methods include:

- CANCEL - To cancel the session being negotiated.
- BYE - To terminate the session, once streaming is completed.
- OPTIONS - To discover a User Agent's response to an invitation without actually signaling the intention (i.e. 'ringing').
- REGISTER - To provide personal mobility.

5.6.2 SIP and User Location

To overcome the limitation of A having to know the terminal address of B in advance, which may be dynamically allocated and forever changing, SIP introduces additional elements to the architecture. These are:

- Proxy Servers.
- Location Servers.
- Registration Servers.
- Redirect Servers.
- Universal Resource Locators (URL).

Every SIP User - including automaton - is given a SIP URL. SIP URLs resemble e-mail addresses, and are of the format: *sip:username@domainname*. Typically, the username is the user's actual name, and the domain Name is the user's home domain (e.g. the ISP) but may also be an independent SIP service provider (similar to the hotmail e-mail service). Within the domain indicated by domain name, there is a SIP Registration Server. Its IP address will be static and easily accessible through DNS (in the same way that mail servers are found when an e-mail is sent to user@domain). The Registration Server listens for messages bearing the REGISTRATION method. Now, when the User Agent starts up, before attempting to start any sessions, the first message it sends is a REGISTRATION. This bears the SIP URL of its user, plus the actual terminal address (IP number), port number, and transport protocol (e.g. TCP, remember that SIP can operate over non-IP networks). Additional optional fields are the time stamp, indicating how long the registration is valid for (the default is one hour), and a preference for being contacted at this location. The Registration Server authenticates the user, and adds the mapping between URL and network address(es) to the Location Server's database.

Figure 3.4 illustrates this.




Figure 5.4: User B registers both his two terminals with a forking SIP proxy server.

SIP URLs allow users to be contacted, irrespective of their current network address. Now, User A simply needs to know the SIP URL of User B, which is constant, as opposed to its possibly ever-changing network address. Knowing a SIP URL is not sufficient to route a message to User B; to do so requires the service of either a SIP Proxy

or Redirect Server. Proxy Servers, as their name suggests, act on User Agents' behalves, routing SIP messages to correct destinations by invoking SIP URL to network address mapping by Location Servers and then forwarding the messages. [Figure 5.5](#) illustrates the revised message flows.

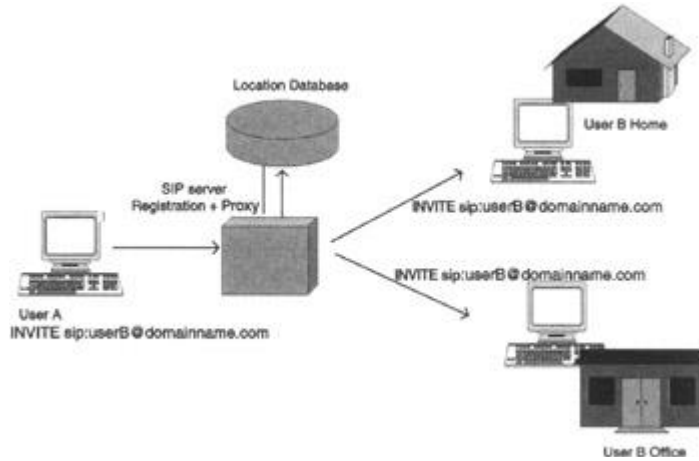


Figure 5.5: User A sends INVITE to user B via proxy server.

User B is currently working from two terminals, each with a User Agent that has registered its network addresses against B's SIP URL. Registrations are additive, although they can be time stamped for periods of validity, and they can be prioritised according to preference in being contacted. When A seeks to contact B, they send their INVITE request to the Proxy, specifying B's URL. The Proxy determines that B currently has two terminal addresses and sends a copy of the message to each, inserting its own address into the path list. B now sends an OK response from one of the terminals to the address at the top of the path list, which results in it being returned to the proxy. The proxy then returns the response to A's User Agent, and remains in the path between A and B for the ensuing ACK. A SIP redirect server is less commonly considered, but acts more like the familiar DNS system. User A would send its INVITE to the SIP server for the domain name (registered with DNS), but the SIP server would return a list of IP addresses to User A, who could then reissue the SIP INVITE direct to User B's terminals.

5.6.3 Characteristics of SIP

- **Simplicity** - SIP has been designed to be very lightweight - it can interoperate with just four headers and three request types. This minimal footprint means that SIP could run on devices with limited processing capabilities - such as pagers or baby alarms. Sessions can be set up in 1.5 round trip times.
- **Generic Session Description** - SIP separates the signaling of sessions from the description of the session. SDP is not mandatory, and SIP could be used to initiate and control completely new types of session.
- **Modularity and extensibility** - SIP is designed to be extensible allowing implementations with different features to be compatible. As will be seen, the UMTS version of SIP is an extension of the basic standard.
- **Programmability** - As will be described in the [next section](#), the introduction of a SIP server offers the possibility of running scripts or code (e.g. Java servlets) that can

alter, re-direct, or copy INVITE or other SIP messages. Not only can SIP servers be used to provide 'Intelligent Network' services like those traditionally seen on voice networks (such as forwarding a call to an answer phone if the phone is busy), but this can be extended to provide intelligent control of advanced multimedia services.

- Integration with other IP component technologies-The design of SIP built heavily on experience of the design of other IP protocols. It is designed to complement IP protocols such as the Real Time Streaming Protocol (RTSP); together, these could be used to offer voice mail services or to invite a video server to play a movie during a multi-party conference.
- Scalability and robustness - SIP servers can be totally stateless, allowing full scalability. There are, however, reasons for having stateful proxies, to provide advanced services, such as those provided by classic call control in 2G networks. SIP also supports multicast sessions, something that is very difficult for traditional circuit based call servers, which require an expensive bridge to connect the parties.

5.7 SIP in Use

5.7.1 Connecting IP and Telephony

Voice is one of the key services that SIP is expected to help support on the Internet it is a real-time peer-to-peer service. However, even in the longer term, it is to be expected that most users world-wide will only have access to the telephone network, and only have voice services. Imagine someone (User A) wants to contact a friend (User B), but User A only has an advanced, fully IP, 3G phone, whereas User B only has a fixed line telephone. How can User B be contacted? What is needed is a gateway - something that sits between two domains that takes in IP voice packets and sends out a PCM 64 kbit/s stream on a PSTN circuit. The gateway also has to take in SIP commands and create SS7 signaling messages (for the PSTN, the SS7 messages are part of a set called ISUP). A SIP PSTN to IP Gateway (SIP PIG) could work as follows.

User A's terminal would create an INVITE message including the E164 (telephone) number of User B, the bit rate and codec(s) that User A had installed on their machine, and their IP address. Within User A's terminal would be a list of SIP proxy servers that provide E164 location services - much like today, all hosts contain a list of default DNS servers to use. User A may simply use a SIP server associated with their UMTS network supplier, but in this case, User B is on a BT network, so User A chooses to send the SIP message to the BT server as this would provide a cheaper service. The SIP proxy server would recognize that User A needed to connect to the PSTN and locate a PIG attached to an appropriate PSTN network. A SIP TRYING message would be returned to User A. User A's INVITE would be forwarded to the PIG, which would in turn seize a circuit-switched trunk termination on the PSTN side and associate it with an RTP termination on the IP side. Once User A received the PIG address, they might then set up some network QoS to the PIG - perhaps with IntServ RSVP messages and when complete, the PIG would select the chosen codec and begin call establishment in the PSTN. The PIG and SIP user agent would exchange messages via the proxy server to signal these events. The PIG sets up the PSTN call with ISUP messages - an Initial Address Message is sent first and the PSTN signals call acceptance with an Address Complete Message. Later, the PSTN sends a Call Progress Message to signal that User B's phone is ringing-this might

be reported back to User A via a SIP RINGING message. For complete details of all the messages exchanged, see the Further reading section. Internally, the PIG must mimic a VoIP client, buffering and decoding the IP packets to create a bit stream this will probably need trans-coding into a 64 kbit/s PCM signal. PIGs are complicated and have many functions: thus, they have been broken down in some VoIP architectures into a media gateway (MG), a Media Gateway Controller (MGC), and a Signaling Gateway (SG), as shown in Figure 5.6. The MG is responsible for all the switching, transcoding, and user-plane aspects. The MGC contains the switch and service functionality.

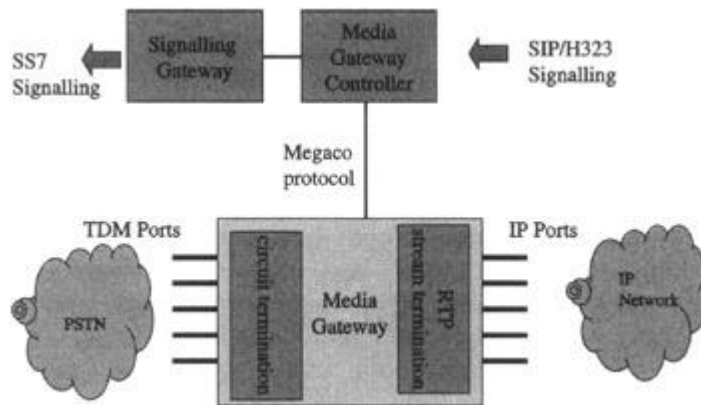


Figure 5.6: PIG in typical VoIP architecture.

The IETF and ITU have jointly standardized the MEGACO (or H.248 in ITU-speak) protocol that is used between the MGC and the MG - the reason for this separation is that MGs might be located remotely from MGCs (the former in exchanges, the latter in server farms, for example). It also allows the two to be separately dimensioned.

5.7.2 SIP Supported Services

SIP has been presented as a major enabler for advanced and multimedia services. This section considers more closely how services such as m-commerce (the mobile version of ecommerce), interactive games, and video applications could be provided using SIP. A number of programming techniques are being developed to allow service creation in SIP networks in general, particularly those involving SIP proxy servers. Thus, some insights can be gained by looking at this topic.

A simple VoIP network using SIP for user location and session negotiation might simply contain a single proxy server, and each PC or mobile terminal would have a User Agent running when they were available to be contacted - so that INVITE messages cause a ringing noise to be generated, for example. The SIP user agents would be interrogated, probably via an API (Application Programming Interface) by the VoIP application - to provide details such as the discovered IP address, or the negotiated codec that the peer VoIP application preferred to use.

If all control messages pass through the SIP proxy server (using a 'VIA server' statement in the SIP header), it is possible to let this hold state and provide services at this point.

For example, users might use a web interface to the SIP proxy server to enable them to set up intelligent call-forwarding, as indicated in [Table 5.1](#).

Table 4.1: Table to indicate call forwarding the preferences of a user			
Calling Party	Time	Handle Call	Priority
Lottery		Current location	Urgent
Mother-in-law		Outer Mongolia tourist information	Non-urgent
Girlfriend	9 a.m.-5 p.m.	E-mail name@domainname.com	

Table 5.1: Table to indicate call forwarding the preferences of a user

There are a number of competing programming methods for creating services at the SIP proxy server:

- CGI scripts - Usual Web scripts that run on Web servers.
- Parlay - A standard telecoms industry interface for IN services.
- JAIN - Java version of Parlay.
- Java servlets - Small java programs that run on the server.
- CPL (Call Programming Language) - A special language with scripts that run on the server.

Each has its own pros and cons - more or less features, security, ease and familiarity of programming, efficiency of operation, and so on. They require state to be kept at the proxy server and also that all the messages related to that session pass through the proxy which SIP can allow. Using this approach of a SIP proxy server holding state, the 3G community has validated that it is relatively easy to recreate the classic IN call services such as call waiting and transfer-on-busy. Unlike IN calls, however, which only work for voice services, these services are independent of the type of application, and so will work for any type of multimedia sessions.

Not only is SIP able to provide the entire set of classic IN services, but this approach can also provide a large range of less common services. These services have proven difficult to provide on traditional IN platforms, despite a clear marketing requirement. A few examples are:

- Third-party call control - A party sets up a call between two other parties without necessarily participating in the call.
- Time-dependent routing - The calls receive different treatments depending on the time of day or the days of week.
- Person-dependent routing - The call is routed to different end points, depending on who is calling. The user might require calls from their boss to be routed to their office desktop, and calls from their family to be routed to the home PC.
- Media-dependent routing - The call is routed to different end points, depending on the type of media requested. The user might prefer, for instance, to receive video on the desktop, instead of the mobile device, where there is only limited bandwidth.
- Calling-name delivery - The name of the caller is displayed on the screen before answering the call.
- Finding a party - As an example, a user willing to play chess can contact the SIP server to request a partner. The INVITE message is addressed to

<sip:chess@bt.com>. The SIP server then makes a look-up in the VHE database, discovers all the users with an interest in chess, and invites them to a session.

Figure 5.7 shows a user registering as interested in local entertainment with their service provider. A content provider, the local theatre, then advertises that 50 low-cost tickets are available. The service provider identifies those most likely to be interested and sets up sessions (for example, an SMS or e-mail), as appropriate.

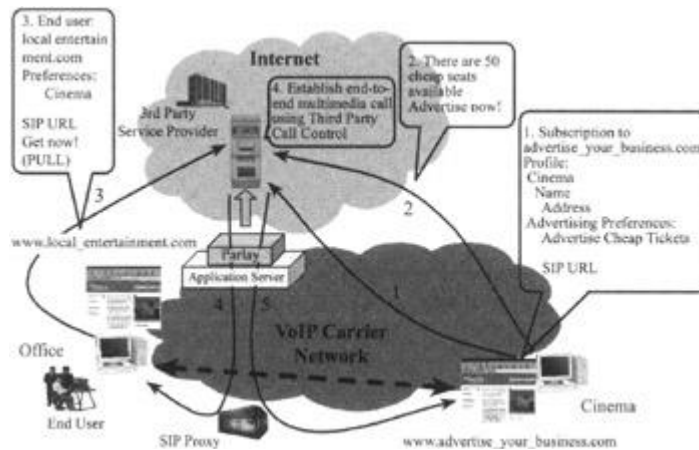


Figure 5.7: Example of SIP service creation.

In reality, certainly in the short term, it is expected that most operators will support standard circuit-switched voice in addition to IP data and multimedia, and also that terminals will be able to use both voice-over-IP and standard telephony. The 3G phone here is a conceptual terminal based on the original 3G vision, and as such has no relationship with a UMTS or CDMA2000 terminal.

5.8 Conclusions

5.8.1 SIP

This chapter began by considering the need for session management for real-time, multimedia applications. SIP was identified as a key protocol to enable users to control the time and manner in which they are contacted. SIP, as common session negotiation protocol, will maximize connectivity for real-time and personal communications. SIP was chosen amongst other contenders because it is a powerful, yet simple and flexible protocol that is likely to play a key role in the future Internet, future UMTS networks, and even in a future IP for 3G network. We presented two examples of the uses of SIP - firstly how SIP can facilitate PSTN Internet inter-working, and secondly how SIP can be used to provide call control services that are terminal and network independent. The rest of this report will touch on other aspects of session control such as the use of RTP to manage a session once established (Chapter 7). SIP itself provides some level of mobility support, in that the location services and SIP renegotiation features allow a user to remain in contact, even if they change terminals during a session (Chapter 6). Although SIP is not in the earliest releases of 3G network standards, the final chapter details how the UMTS community is considering utilizing SIP in the near future.

In addition to these roles, the session initiation protocols can be used in more advanced ways. For example, a network server that assists in session initiation could interpret the session descriptions and then act as a bandwidth broker to install the required QoS information into the network. However, this level of integration is not assumed to be in accordance with the Internet principles and may, from the end user's perspective, have security implications.

5.8.2 VHE

SIP has been claimed as a key element in delivering the VHE concept. The VHE concept is about:

- A single bill.
- A single number.
- Common operating and call control procedures.
- A place to store user preferences and data.
- Something to tailor a service to the connection and terminal being used.

Within this report, the operator-specific and commercially sensitive issue of billing is avoided. In a model where users can be contacted only through a SIP proxy server, it is possible to see that the SIP server could also act as a coordination point for all billing activity.

SIP servers do not provide a single number for a user - they provide something much more attractive a single name for a user. This can be achieved either through the use of a full proxy server or simply through the use of a redirect server with access to the location server. This single name can be used for video as well as voice services. Well-established mail systems will probably still retain their independence, and as such their own naming schemes. They are store and forward systems, which means that a message can be sent even when the intended recipient is not on any network. SIP is basically aimed at supporting instant communications. However, as indicated above, SIP proxies could be used to tell a calling party that the only type of communication that the recipient is prepared to accept is an e-mail.

SIP is an open, simple standard. It is totally independent of the network over which it operates. Thus, users of SIP will have the benefits, for example, of easy individualized services, which will be available to the user independently of the network - thus, these services will function correctly, even when a user roams from their home network. These are the goals of having common operating and call control procedures.

SIP allows user data and preferences to be stored either in a user's own terminal or in a proxy server. The advantage of the proxy server is that a user can move between terminals, for example when they need to recharge the battery on the mobile.

Finally, SIP is fundamentally about enabling the characteristics of a session to be tailored to the terminal and network through which a user is connected. This is the basic functionality of SIP - the ability to negotiate the type of service that will be used.

Thus, SIP can be seen to provide the full VHE vision. However, it is worth remembering that it is not the only way to achieve this vision. For example, 2G operators are also continually developing their networks in order to support such services. The CAMEL (Customized Applications for Mobile network Enhanced Logic) platform is being developed for this purpose. This enables 2G operators to offer services, which can still be accessed whilst a user is roaming away from their home network. However, CAMEL is limited. It only supports circuit-switched voice services (such as short code dialing) and has no mobility support. Thus, a user could not switch terminals, or insist that a certain acquaintance only e-mails them while at work.

From a user's perspective, SIP has a further advantage over the 2G approach to advanced service provision: it is much easier to separate network connectivity from the session management functionality. Indeed, SIP can run without any co-operation from any network components. Today, people choose to join a specific network partly because of the services it offers. With SIP, there is no reason why a user could not add the SIP functionality themselves. If a user wanted more than basic session negotiation, they would simply use their home PC that was 'always on', register a domain name, and start a shareware SIP proxy or re-direct server on it. The user could then tell their friends their new name, and obtain advanced services, at no additional cost. They could then change their operator without needing to reinstall all their preferences, or change their SIP address. A server could then be run from home as a small business. Whilst some network operators, certainly within the UK, are looking to avoid people operating servers at home, certainly they cannot prevent a small business providing this service. This bypasses a potential source of operator 'lock-in'. Indeed, users may be able to be registered with different names with different SIP providers, for example a business address and a home address, yet use one network and one terminal.

Chapter 6

IP Mobility

6.1 Scope

This chapter will provide an overview of IP mobility. It aims to be pretty self-contained, and so should stand alone fairly independently of the other chapters.

IP mobility is very important, because it is predicted that the vast majority of terminals will be mobile in a few years and that the vast majority of traffic will originate from IP-based applications. The challenge of 'IP mobility' is to deliver IP-based applications to mobile terminals/users, even though, traditionally, IP-protocols have been designed with the assumption that they are stationary.

In outline, this chapter considers:

- The distinction between personal and terminal mobility, and between an identifier and a locator.
- For terminal mobility the distinction between macro (or global) and micro (or local) mobility.
- Tunnel-based and per-host forwarding approaches to micromobility - Their key
- Features and how they compare.
- Other aspects of terminal mobility - Context (or state) transfer, paging, and security.

As part of this, the chapter includes an outline of various protocols:

- SIP (Session Initiation Protocol) - Its use for personal and macromobility.
- Mobile IP - For macromobility.
- Hierarchical mobile IPv6, regional registration, fast mobile IP for v4 and v6, cellular
- IP for v4 and v6, Hawaii, MER-TORA - For micromobility.

The chapter does not consider MANETs (mobile ad hoc networks): networks without a fixed infrastructure . In other words, the chapter concentrates on how to cope with mobility in an IP network reminiscent of a traditional cellular network - that is, a fixed network with base stations that provide wireless connections to mobile terminals.

The treatment is at quite a high level; the aim is to provide an introduction to the subject, to enable the reader to understand what the key issues are, and hopefully to help an incisive analysis of future proposals. The chapter also aims to give a flavour of some of the latest thinking on this fast moving subject.

6.2 Introduction - What is IP Mobility?

This part covers a number of topics that explore what is meant by 'IP mobility'. First, two (complementary) types of mobility are distinguished: personal and terminal. Second, the different protocol layers that mobility can be solved at are looked at. Third, we discuss how the distinction between an identifier and a locator offers an insight into mobility.

6.2.1 Personal and Terminal Mobility

A traditional mobile network like GSM supports two types of mobility: terminal and personal.

Terminal mobility refers to a mobile device changing its point of attachment to the network. The aim is that during a session, a mobile terminal can move around the network without disrupting the service. This is the most obvious feature that a mobile network must support.

Personal mobility refers to a user moving to a different terminal and remaining in contact. 2G networks have a form of personal mobility, because a user can remove their SIM card and put it in another terminal - so they can still receive calls, they still get billed, and their personal preferences like short dialing codes still work.

What mobility is widely available in the Internet today? First, portability, which is similar to terminal mobility, but there is no attempt to maintain a continuous session. It deals with the case where the device plugs into a new network access point in between sessions. For example, a user can plug in their laptop into any network port on their home network, for example the one which happens to be nearest to where they are working. However, true terminal mobility is not currently widely available in the Internet today. Second, personal mobility, for example through a WWW portal (such as Yahoo), enables users to send and receive web-based e-mail from Internet cafes. However, this type of solution is limited in that it only operates through the portal.

The bulk of this chapter considers various techniques and protocols that would enable IP terminal mobility. [Section 6.3](#) also briefly considers how an IP network can effectively support personal mobility.

6.2.2 The Problem of IP Mobility

Broadly speaking, there are three ways of viewing the 'problem of IP mobility', corresponding to the three layers of the protocol stack that people think it should be solved at:

- Solve the problem at Layer 2 - This view holds that the problem is one of 'mobility' to be solved by a specialist Layer 2 protocol, and that the movement should be hidden from the IP layer.
- Solve the problem at the 'application-layer' - This view similarly holds that IP layer should not be affected by the mobility, but instead solves the problem above the IP layer.
- Solve the problem at the IP layer - Roughly speaking, this view holds that 'IP mobility' is a new problem that requires a specialist solution at the IP layer.

6.2.3 Layer 2 Solutions

This approach says that mobility should be solved by a specialist Layer 2 protocol. As far as the IP network is concerned, mobility is invisible - IP packets are delivered to a router and mobility occurs within the subnet below. The protocol maintains a dynamic mapping between the mobile's fixed IP address and its variable Layer 2 address and is equivalent to a specialist version of Ethernet's ARP (Address Resolution Protocol). This is approach

taken by wireless local area networks (LANs), e.g. through the inter-access point protocol (IAPP). Although such protocols can be fast, they do not scale to large numbers of terminals. Also, a Layer 2 mobility solution is specific for a particular Layer 2, and so inter-technology handovers will be hard.

Another example is where a GSM user dials into their ISP, with PPP used to give an application level connectivity to their e-mail or the Internet. Mobility is handled entirely by the GSM protocol suite and IP stops at the ISP - so as far as IP is concerned, the GSM network looks like a Layer 2. Clearly, this solution does work and indeed has been very successful. However, the problem is that all the IP protocols must be treated as applications running from the mobile to the ISP. The implication is that many IP protocols cannot be implemented as intended - for example, it is not possible to implement web caching or multicasting efficiently. These protocols will become increasingly important in order to build large efficient networks.

6.2.4 Application-layer Solutions

Although generally called application-layer solutions, really this term means any solution above the IP layer. An example here would be to reuse DNS (Domain Name System). Today, DNS is typically used to resolve a website's name (e.g. www.bt.com) into an address (62.7.244.127), which tells the client where the server is with the required web page. At first sight, this is promising for mobility, and in particular for personal mobility; as the mobile moves, it could acquire a new IP address and update its DNS entry, and so it could still be reached. However, DNS has been designed under the assumption that servers move only very rarely, so to improve efficiency, the name-to-address mapping is cached throughout the network and indeed in a client's machine. This means that if DNS is used to solve the mobility problem, often an out-of-date cached address will be looked up. Although there have been attempts to modify DNS to make it more dynamic, essentially by forcing all caching life times to be zero, this makes everyone suffer the same penalty even when it is not necessary.

[Section 6.3](#) examines another IP protocol, SIP, for application-layer mobility.

6.2.5 Layer 3 Solutions

The two previous alternatives have limited applicability, so the IP community has been searching for a specialist IP-mobility solution, in other words, a Layer 3 solution. It also fits in with one of the Internet's design principles: 'obey the layer model'. Since the IP layer is about delivering packets, then from a purist point of view, the IP layer is the correct place to handle mobility. From a practical point of view, it should then mean that other Internet protocols will work correctly. For example, the transport and higher-level connections are maintained when the mobile changes location. Overall, this suggests that Layer 3 and sometimes Layer 2 solutions are suitable for terminal mobility, and 'application' layer solutions are sometimes suitable for personal mobility.

6.2.6 Locators vs. Identifiers

One way of thinking about the problem of mobility is that we must have some sort of a dynamic mapping between a fixed identifier (*who* is the mobile to whom packets are to be delivered?) and a variable locator (*where* in the network are they?). So, for instance, in

the DNS case, the domain name is the identifier, and the IP address is the locator. Similarly, the www portal (e.g. Yahoo) would have the user's e-mail address (for example) as their identifier and again their current IP address as the locator ([Table 6.1](#)).

Since mobility is so closely tied to the concept of an identifier, it is worth thinking about the various types of identifier that are likely:

- **Terminal ID** - This is the (fixed) hardware address of the network interface card. A terminal may actually have several cards.
- **Subscription ID** - This is something that a service provider uses as its own internal reference, for instance so that it can keep records for billing purposes. The service provided could be at the application or network layer.
- **User ID** - This identifies the person and clearly is central to personal mobility. During call set-up, there could be some process to check the user's identity (perhaps entering a password) that might trigger association with a subscription id. In general, a user ID might be associated with one or many subscription ids, or vice versa.
- **Session ID** - This identifies a particular voice-over-IP call, instant messaging session, HTTP session, and so on. Whereas the other three ids are fixed (or at least long-lasting), the session ID is not.

So, personal mobility is really about maintaining a mapping between a user ID and its current terminal ID(s), whereas terminal mobility is about maintaining the same session ID as the terminal moves.

What is the role of an IP address? From the perspective of an IP network, the main role of an IP address is to act as a locator, i.e. it is the piece of information that informs the IP routing protocol where the end system is (or, to put it more accurately, it allows each router, on a hop by-hop basis, to work out how to direct packets towards the end system). A change of location therefore implies a change of IP address.

However, a typical application today also uses the IP address as part of the session identifier. This does not cause a problem in the fixed Internet-even if the terminal gets allocated IP address(es) dynamically. For instance each time it is re-booted through DHCP, the new voiceover-IP call (or whatever) will simply use the new IP address. But if the terminal is mobile, we have a conflict of interest: the IP address is acting as both an identifier and a locator -implying that the IP address should be both kept and changed. This 'functionality overload' is the real problem that IP terminal mobility solutions tackle. The two main approaches are:

- To allocate two IP addresses to the mobile - one of which stays constant (the identifier) and one of which varies (the locator). This approach is said to be tunnel based or mobile IP-based.
- To have one IP address (the identifier) plus a new routing protocol (which handles the variable location). This approach is called per-host forwarding.

Some other relevant ideas are:

- To re-write applications so that they can support a change in IP address - for example, the restart facility in some versions of FTP. This is called 'application-layer recovery'.
- Similarly, to re-write the transport protocols so that they can support a change in IP address (e.g. through a new TCP option that allows a TCP connection to be identified by a constant 'token', which maps to the changing IP address).
- To invent a new 'Host Identity'. Transport connections would be bound to the host Identity instead of the IP address. This approach is at an early stage of exploration at the IETF.

An (open) question is whether these ideas would allow for 'seamless handovers', i.e. no noticeable degradation in quality of service during the handover. They might be better considered as approaches for making portability better, or as things that complement terminal mobility.

Incidentally, this (correctly) suggests that one of the hardest problems to deal with is a mobile server. Luckily, these are very rare today, but it is possible that they could be common one day (maybe mobile webcams). The problem is not considered further here.

There is also a terminological overload: a 'locator' is often called an 'address'. This can cause some confusion, since an 'IP address' is an 'identifier' as well as a 'locator'.

In any case, application-layer recovery is a good feature in wireless environments, because the link to the mobile may go down.

6.3 SIP - A Protocol for Personal Mobility

The basic operation and primary usage of SIP, the Session Initiation Protocol, is described in [Chapter 5](#). This section briefly considers how SIP can be used to provide personal mobility. Essentially, SIP supports a binding between a user-level identifier (the SIP URL) and the user's location, which is the name of the device where the user can be currently found. SIP can provide such personal mobility either at the set-up of a call or during the media session.

- **At Call Set-up** - At present User A must use a different name or number to contact User B, depending on whether User A wants to talk on the phone, send an e-mail, engage in an instant messaging session, and so on. SIP enables User B to be reached at any device via the same name (sip: <phil@abctel.com>). When User A wants to contact User B, User A's SIP INVITE message is sent to User B's SIP proxy server, which queries the location database (or registrar) and then sends the INVITE on to one of User B's devices, or alternatively 'forks' it to several, depending on User B's preferences. User B can then reply (SIP OK) from the device that they want to use.

See [Figure 6.4](#) in [Chapter 6](#). User B could also advertise different SIP addresses for different purposes, for example work and personal - just as with e-mail today. This might allow User B's SIP server to make a more intelligent decision about how to deal with an INVITE.

- During a Media Session - This sits somewhere between personal and terminal mobility and refers to the ability of a user to maintain a session whilst changing terminals. It is sometimes called service mobility. For example, User A might want to transfer a call that started on their mobile phone on to the PC when they reach the office, or they might want to transfer the video part of a call on to a high-quality projector. The main SIP technique to achieve such session mobility is to explicitly transfer the session to the new destination using the REFER request message - see Figure 6.1. The REFER tells User B to re-INVITE User A at User A's address ; the call-ID is included so User A knows that this is not a fresh INVITE. Alternatively, User A could send the REFER to their new terminal, and it would then send the re-INVITE to User B.

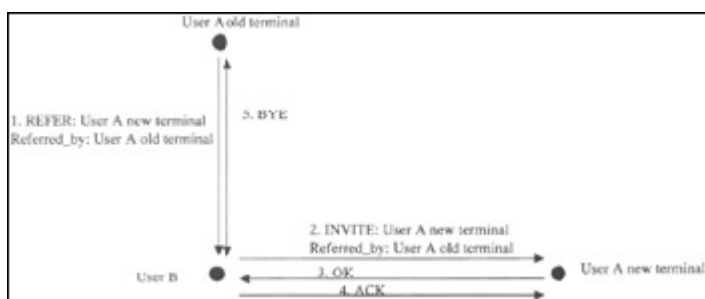


Figure 6.1: Use of SIP REFER message for application-layer mobility (User A moves on to a new terminal).

This implies that the application must be able to cope with application-layer recovery.

6.4 Introduction to Terminal Mobility

The rest of this chapter considers terminal mobility in an IP network, covering the 'IP layer' solutions. This section briefly considers the important distinction between (terminal) macro and micro mobility. Subsequent sections look at some specific approaches and protocols for macro mobility (in particular Mobile IP, but also briefly the possible use of SIP for terminal mobility) and micro mobility (in particular, the tunnel-based protocols of hierarchical mobile IP and fast mobile IP, and the per-host forwarding protocols of cellular IP, Hawaii and MERTORA). The chapter then compares the various micro mobility protocols. Finally, it looks at some other features that are important for a complete terminal mobility solution (paging, context transfer, and security).

The basic job of a terminal mobility protocol is to ensure that packets continue to be delivered to the mobile terminal, despite its movement resulting in it being connected through a different router on to the network. The main requirements are that the protocol does this:

- Effectively - Including for real time sessions.
- Scalably - For big networks with lots of mobiles.
- Robustly - For example to cope with the loss of messages.

6.4.1 Macromobility vs. Micromobility

It is generally agreed that IP terminal mobility can be broken into two complementary parts - macromobility and micromobility - and that these need two different solutions. These terms are generally used informally to mean simply 'mobility over a large area' and 'mobility over a small area'. It might seem a little strange that such woolly definitions should lead to such firm agreement that there needs to be two different solutions. In fact, the important distinction is between terminal mobility to a new administrative domain (AD) and within the same AD. For example, a mobile might move around a campus wireless network, handing over from one wireless LAN base station to another, and then off on to a public mobile network. These handover cases are significantly different, because an inter-AD handover implies that:

- The mobile host needs to be re-authenticated, because the security/trust relationship is much weaker between ADs than within one.
- The user's charging regime, priority, and QoS policy are all likely to be changed.
- A different IP address must be used (because IP addresses are owned by the AD), whereas it may or may not be for an intra-AD handover (it depends on the particular micromobility protocol).
- Issues such as the speed and performance of the handover are less relevant, simply because such handovers will be much rarer.
- There is no guarantee of mobility support in the new AD, because the protocols being run are not certain, and therefore an inter-AD handover must rely on protocols that can exist outside the two ADs involved.

It is thus suggested that two complementary protocols are needed: one solving the macromobility problem and one the micromobility.

However, as will be discussed later, micromobility protocols implicitly assume mobility within an Access Network (rather than within an Administrative Domain). The terminology used is (see also [Figure 6.2](#)):

- An Access Network (AN) is simply a network with a number of Access Routers, Gateway(s) and other routers.
- An Access Router (AR) is the router to which the mobile is connected, i.e. that at the 'edge' of the Access Network. It is an IP base station.
- An Access Network's Gateway (ANG) is what connects it to the wider Internet.
- The other Routers could be standard devices or have extra functionality to support IP micromobility or quality of service.




Figure 6.2: Terminology for Access Network.

The Access Network and Administrative Domain may correspond to each other, but they may not; for example, the operator could design an AN on technical grounds (e.g. how well does the micromobility protocol scale?), rather than the commercial focus of the AD (e.g. interworking agreements with other operators). This leaves a 'hole', i.e. an 'inter-AN, intra-AD handover'; at present, it seems that a macromobility protocol is fully adequate to handle this.

Finally, on a terminological point, some people do not like the term 'micromobility', basically because it has been used to mean a variety of slightly different things over the years, and so can cause confusion. Alternative terms include intra-access network mobility, localized mobility management and local mobility. Also, an alternative term for 'macromobility' is global mobility.

Being used in the sense [draft-ietf-mobileip-reg-tunnel-03.txt] Domain: A collection of networks sharing a common network administration.

6.5 Mobile IP - A Solution for Terminal Macromobility

6.5.1 Outline of Mobile IP

The best-known proposal for handling macromobility handovers is Mobile IP. Mobile IP has been developed over several years at the IETF, initially for IPv4 and now for IPv6 as well. Mobile IP is the nearest thing to an agreed standard in IP-mobility. However, despite being in existence for many years and being conceived as a short-term solution, it still has very limited commercial deployment (the reasons for this are discussed later); Mobile IP products are available from Nextel and ipUnplugged, for example.

In Mobile IP, a mobile host is always identified by its home address, regardless of its current point of attachment to the Internet. Whilst situated away from its home, a mobile also has another address, called a 'Care-of Address' (CoA), which is associated with the mobile's current location. Mobile IP solves the mobility problem by storing a dynamic mapping between the home IP address, which acts as its permanent identifier, and the care-of address, which acts as its temporary locator.

The key functional entity in mobile IP is the Home Agent, which is a specialised router that maintains the mapping between a mobile's home and care-of addresses. Each time the mobile moves on to a new subnet (typically, this means it is moved on to a new Access Router), it obtains a new CoA and registers it with the Home Agent. Mobile IP means that a correspondent host can always send packets to the mobile: the correspondent addresses them to the mobile's home address - so the packets are routed to the home link-where the home agent intercepts them and uses IP-in-IP encapsulation (usually) to tunnel them to the mobile's

CoA. (In other words, the home agent creates a new packet, with the new header containing the CoA and the new data part consisting of the complete original packet, i.e. including the original header.) At the other end of the tunnel, the original packet can be extracted by removing the outer IP header (which is called decapsulation). (Figure 6.3a and 6.3b).

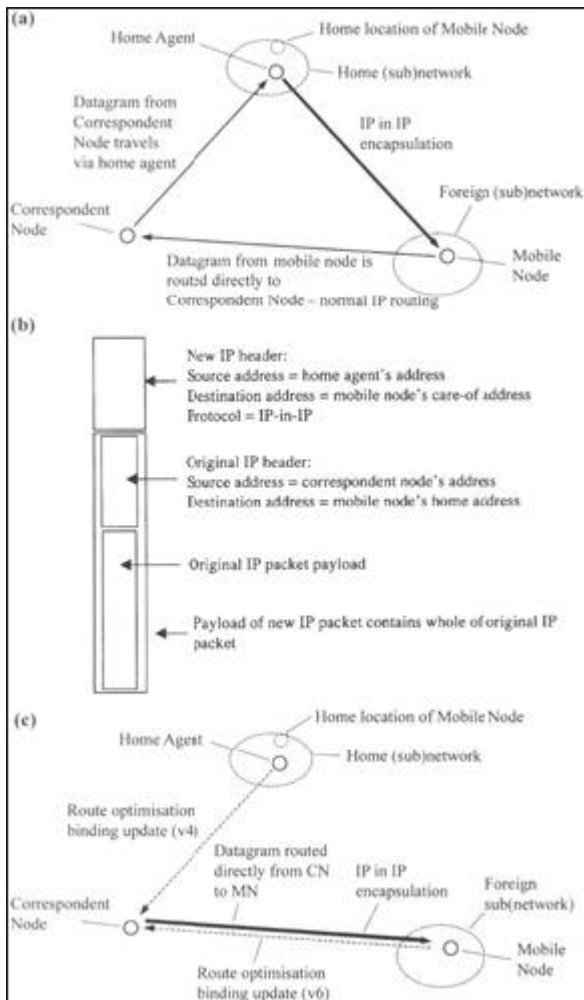


Figure 6.3: Mobile IP. (a) Triangular registration and routing (b) Packet encapsulation (c) Route optimisation.

Note that mobile IP is only concerned with traffic to the mobile - in the reverse direction, packets are sent directly to the correspondent host, which is assumed to be at home. (If it is not, mobile IP must be used in that direction as well.)

A couple of key features of mobile IP are:

- It is transparent to applications. They can continue to use the same IP address, because the home agent transparently routes them to the mobile's current care-of address.
- It is transparent to the network. The network's standard routing protocol continues to be used. Only the mobiles and the home agent (and foreign agents - see later) know about the introduction of mobile IP - other routers are unaffected by it.

6.5.2 Mobile IPv4

The Mobile IPv4 protocol is designed to provide mobility support in an IPv4 network. As well as the Home Agent (HA), it introduces another specialised router, the Foreign Agent (FA). For example, each access router could be a FA. A mobile node (MN) can tell which FA it is 'on' by listening to 'agent advertisements', which are periodically broadcast by each FA.

The advertisement includes the FA's network prefix. When the MN moves, it will not realize that it has done so until the next time it hears a FA advertisement; it then sends a registration request message. Alternatively, the MN can ask that an agent sends its advertisement immediately, instead of waiting for the periodic advertisement.

Mobile IPv4 comes in two variants, depending on the form of its CoA. In the first, the MN uses the FA's address as its CoA and the FA registers this 'foreign agent care-of address' (FACoA) with the HA. Hence, packets are tunnelled from the HA to the FA, where the FA decapsulates and forwards the original packets directly to the MN. In the second variant, the MN obtains a CoA for itself, e.g. through DHCP, and registers this 'co-located CoA' (CCoA) either directly with the HA or via the FA. Tunnelled packets from the HA are decapsulated by the MN itself.

The main benefit of the FA-CoA approach is that fewer globally routable IPv4 addresses are needed, since many MHs can be registered at the same FA. Since IPv4 addresses are scarce, it is generally preferred. The approach also removes the overhead of encapsulation over the radio link, although, in practice, header compression can be used to shrink the header in either the FA-CoA or CCoA scenario.

There are several problems with Mobile IPv4, which can be alleviated with varying success. These are discussed below.

6.5.3 Triangular Routing and Route Optimization

In the basic Mobile IPv4 described above, all packets from the correspondent node (CN) go via the HA to the MN. This 'triangular' route can be very inefficient - imagine a visitor from Australia to England communicating with someone in the same office. An optional extension to MIP, called Route Optimization allows a CN to send packets directly to a MN. It works by the HA sending a binding update to the CN, in response to mobile node warnings or correspondent node requests. (Figure 6.3c). However, route optimization does require an update to the CN's protocol stack (so it can cache the MN's CoA and do encapsulation), and it may not be useful in some circumstances (e.g. if the MN has signed up to many servers that 'push' information occasionally).

6.5.4 Reverse Tunneling

Mobile IPv4 suffers from a practical problem with firewalls (or, more generally, a router that performs ingress filtering). A MN uses its home address as its source address, but a firewall expects all packets within its network to use a topologically correct source address (i.e. to use the same network prefix) and will therefore throw away packets from the MN. To circumvent this, an extension has been added, known as Reverse Tunnelling. It establishes a 'reverse tunnel', i.e. from the care-of address to the home agent. Sent

packets are then decapsulated at the home agent and delivered to correspondent nodes with the home address as the IP source address.

6.5.5 NAT Traversal

Similarly, Mobile IPv4 suffers from a practical problem with Network Address Translators (NATs). NATs are discussed in more detail in [Chapter 5](#). They are used extensively in IPv4 networks, owing to the shortage of publicly routable IPv4 addresses. They allow many IP nodes 'behind' a NAT to share only a few public addresses, and indeed for several nodes to share the same address simultaneously, whilst using different port numbers. The latter is particularly problematic for Mobile IP: the HA (or CN) tunnels packets, using IP-in-IP encapsulation, to the MN's publicly routable care-of address; when the packets reach the NAT, it must translate the address to the MN's actual private care-of address - but if several MNs are sharing the same address, it cannot do this. A proposed solution involves using IPin-UDP encapsulation; the UDP header carries the extra information about the port number, which allows the NAT to identify the correct MN.

6.5.6 Address shortage

Even when FA-CoAs are used, the MN still needs a home address. The shortage of IPv4 addresses means that an ISP or network operator would much rather give each user an address dynamically (through DHCP).

6.5.7 Foreign Agents

The need to deploy FAs has perhaps proved the biggest stumbling block to the deployment of Mobile IPv4: It is extra kit for a network operator to buy; the mobile loses service if it moves on to a network without foreign agents; it makes security harder to implement, because the home agent must trust the foreign agents; and it is in tension with the end-to-end IP design principle, because there is a point in the network that modifies the packet.

6.5.8 Mobile IPv6

Mobile IPv6 is designed to provide mobility support in an IPv6 network. It is very similar to Mobile IPv4 but takes advantage of various improved features of IPv6 to alleviate (solve) some of Mobile IPv4's problems.

- Only CCoAs need to be used, because of the increased number of IPv6 addresses.
- There are no foreign agents. This is enabled by the enhanced features of IPv6, such as Neighbour Discovery, Address Auto-configuration, and the ability of any router to send Router Advertisements.
- Route Optimization is now built in as a fundamental (compulsory) part of the protocol. Route Optimization binding updates are sent to CNs by the MN (rather than by the home agent).
- There is no need for reverse tunnelling. The MN's home address is carried in a packet in the Home Address destination option. This allows a MN to use its care-of address as the Source Address in the IP header of packets it sends - and so packets pass normally through firewalls.

- Packets are not encapsulated, because the MN's CoA is carried by the Routing Header option added on to the original packet. This adds less overhead costs and possibly simplifies QoS (see later).
- There is no need for separate control packets, because the Destination Option allows control messages to be piggybacked on to any IPv6 packet.

6.5.9 Relationship of SIP and Mobile IP

Earlier, the use of the Session Initiation Protocol (SIP) for personal (including session) mobility was described. However, SIP can also be used for terminal macromobility. The idea is conceptually very similar to mobile IP.

A mobile node re-registers with its SIP location database each time it obtains a new IP address - this is just like the binding updates to the home agent in mobile IP. A correspondent wishing to communicate with the MN sends a SIP INVITE, which reaches the MN's SIP server. If this is a SIP proxy server, it forwards the INVITE to the MN at its current IP address, whereas if it is a SIP redirect server, it tells the correspondent the MN's IP address so that I can ask directly. This is reminiscent of the versions of mobile IP without and with route optimisation, respectively.

If the MN moves during a call, it can send the correspondent another INVITE request (with the same call identifier) with the new address (in the CONTACT field and inside the updated session description). This is very similar to the session mobility described earlier.

So, is there any difference between SIP and mobile IP for terminal mobility? Well, whereas mobile IP requires the installation of home agents and modifications to the mobile's operating system (and the correspondents if route optimisation is used), SIP requires the presence of SIP servers and that the host and correspondent run the SIP protocol. So, in some ways, the question of whether SIP or mobile IP is better for terminal mobility is really a judgement about which protocol will turn out to be more successful. Favouring SIP is its wide functionality and its use in the IMS (Internet Multimedia Subsystem) of UMTS Release 5, whereas Mobile IP's backers could point to its longevity and use in 3GPP2, for example.

Of course, it is quite possible to believe that both SIP and mobile IP will have a role and that actually they will complement each other. There are a number of ways in which this could happen. For example, SIP could be used for personal mobility and mobile IP for terminal macromobility, by registering the home address with the SIP server; as variants, the SIP server could use the home agent as its location register, or the mobile could register its CoA with the SIP server. Another option is for macromobility to be supported by mobile IP for long-lived TCP connections (e.g. FTP), and by SIP for real-time sessions.

A header option means that the normal IP packet header is extended with an optional field carrying useful information..

In fact, packets sent via the Home Agent, i.e. before Route Optimization, cannot use the Routing Header without compromising security, and so the HA must tunnel packets to the MN's CoA.

In fact, the requirement is slightly stronger than this for TCP applications, where the TCP connection needs to be maintained during a move. One possible solution is that a mobile uses a TCP tracking agent (SIP-EYE) to maintain a record of ongoing TCP connections, and when it hands over, it sends a SIP INFO message to the correspondent asking for the mobile's old address to be bound to its new address. This is very reminiscent of route-optimised mobile IP with co-located care-of addresses.

6.6 Terminal Micromobility

6.6.1 Introduction

This is quite a long section, and the reader is encouraged to skip between areas of particular interest. One reading route is to tackle the introduction sections (this one, plus the introductions to local mobility agent schemes, fast and smooth mobile IP schemes, and perhost forwarding protocols), perhaps followed by the comparison section or the specifics of a particular protocol.

The obvious way to provide (terminal) micromobility is simply to use mobile IP. However, this presents a number of problems, some of which are:

- Handovers may be slow, because the mobile must signal its change of care-of address (CoA) to the home agent. This may take a long time if the home agent is far away, perhaps in a different country.
- The messaging overhead may be significant, particularly if the home agent is distant, as this will induce signaling load in the core of the Internet.
- Mobile IP may interact with quality of service (QoS) protocols, thus making QoS implementation problematic. For example, mobile IP utilises tunnels, and so packet headers - which may contain QoS information - become invisible.

Instead, researchers suggest that a more specialised protocol is needed to deal with micromobility. We will assume in the discussion below that the packets 'somehow' have been delivered to an access network's (AN's) gateway, or else that they have originated within the AN (i.e. a mobile to mobile call).

There has been a huge amount of work on the micromobility problem, with many different ideas and protocols suggested.

Broadly speaking, there are two ways of dealing with micromobility:

- Mobile IP-based schemes - these extend basic mobile IP. They are characterised by the use of tunnelling (and Router headers in IPv6), and in general by the mobile acquiring a new care-of address (CoA) each time it moves.
- Per-host forwarding - these introduce a dynamic Layer 3 routing protocol in the AN. In general, the mobile keeps its CoA whilst it remains in the AN.

There are two common aims to improve on basic mobile IP:

- To reduce the signaling load by localising the path update messages to within the AN or some part of it. This is done by introducing mobility functionality on one, or some, or all of the routers in the access network so that the Home Agent can remain unaware that the MH has moved.

- To speed up handovers, so, from a mobile's point of view, its application does not see a significant delay and suffers no loss of packets. Such handovers are, respectively, said to be 'fast' and 'smooth', or 'seamless' if both apply.

6.7 Mobile IP-based schemes

Two complementary threads of work have been taking place.

6.7.1 Local Mobility Agents

These have been developed on the basis that mobile IP is *almost* the right way of doing it. They assume that mobile IP's problems arise only from the potentially long distance signaling back to the home agent when a mobile moves, which can be solved by introducing a local proxy mobility agent. In this way, when the mobile changes its CoA, the registration request (usually) does not have to travel up to the home agent but remains 'regionalised'. These schemes are predominantly concerned with reducing the signaling load, compared with basic mobile IP.

6.7.2 'Fast and Smooth' Mobile IP-based Schemes

This refers to a variety of 'tricks' introduced to try to make the mobile IP handover seamless (reduction of signaling is not particularly a concern). The most important idea is to use supplementary information to work out that a handover is probably imminent (for instance, this could be link layer power measurements) and to take proactive action on the mobile's behalf. The main steps are to acquire a new CoA that the mobile can use as soon as it moves on to the new access router (AR), and to build a temporary tunnel between the old and new ARs, which stops any packets being lost whilst the binding update messages are being sent.

6.7.3 Per-host Forwarding Schemes

In per-host schemes, the information about the location of the mobiles is spread across several of the routers in the access network. In terms of the earlier discussion, the mapping between a mobile's identifier and its locator is distributed rather than centralised. The location information simply indicates the next router to forward a packet on to, rather than its final destination. Compared with basic mobile IP, these schemes are generally concerned with both reducing the signaling load and speeding up handovers. Three broad techniques for per-host forwarding have been explored:

6.7.4 New Schemes

These schemes assume that it is best to design a new, dynamic Layer 3 routing protocol to operate in the AN. The new protocol installs forwarding entries for each Mobile Host (MH) within the Access Network. The well-known Cellular IP and HAWAII (Handoff-Aware Wireless Access Internet Infrastructure) protocols fall into this category.

6.7.5 MANET-based Schemes

MANET protocols were originally designed for Mobile Ad hoc NETWORKs, where both hosts and routers are mobile, i.e. there is no fixed infrastructure and the network's topology changes often. Clearly, therefore, a MANET protocol can cope with our scenario, where there is a fixed infrastructure, and only hosts can be mobile, although one would expect some modifications to optimise the protocol.

6.7.6 Multicast-based Schemes

The claim here is that the mobility problem is rather like the multicast problem, in that in neither case is a terminal in a fixed, known place. The basic idea is that the protocol builds a multicast 'cloud' centred on the MH's current location but which may also cover where it is about to move to.

Typically, per-host forwarding schemes have the following characteristics:

- The way in which IP addresses are assigned is unrelated to the mobile's current
- position within the network topology. This is substantially different from IP address assignment in the normal Internet.
- There is no encapsulation or decapsulation of packets. Amongst other things, this avoids the overhead associated with mobile IP-based schemes.
- Signalling is introduced to update the mobile specific routes, which is interpreted by several routers within the access network (whereas signaling in mobile IP-based protocols is transmitted transparently by the AN's routers between the mobiles and the mobility agents).

Figure 6.4 shows a family tree for some IP mobility protocols. The references provide further details on the various protocols discussed.

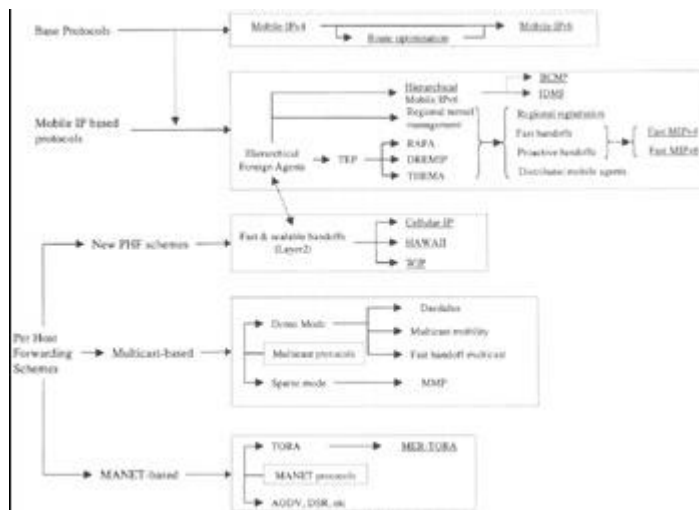


Figure 6.4: Protocol family tree for IP mobility (underlined protocols are discussed in this Chapter).

6.8 Mobile IP-based Protocols

6.8.1 Local Mobility Agent Schemes

These protocols introduce a local mobility agent, which is just a specialised router that essentially acts as a local proxy for the home agent. When a mobile moves, it normally hands over between two access routers that are 'under' the same local mobility agent. Hence, it only needs to inform this mobility agent; the Home Agent and correspondent hosts remain unaware of the move. There are two things this should achieve:

- Reduce the amount of signaling - there are fewer messages (just one local update, rather than one to the home agent and - assuming route optimisation - one to each correspondent) and also a shorter distance for the messages to travel.
- Reduce the latency associated with a handover - because the update only has to travel as far as the local mobility agent. So, the users will see a shorter break in their communications.

The basic method is that as well as the standard home address and care-of address, the mobile has an extra care-of address (CoA) that is associated with the local mobility agent. The home agent remembers which local mobility agent the mobile is on, whereas the local mobility agent can tunnel packets on towards the correct AR. A correspondent host sends its packet either to the home agent or to the local mobility agent after route optimisation. There is no attempt to route-optimize further, i.e. to the CoA associated with the current AR, since this would involve extra signaling and degrade (at least part of) the advantage of the local mobility agent.

It is also suggested that there can be a hierarchy of local mobility agents, so that packets would be sequentially tunnelled from one to the next. However, this is generally opposed, owing to the processing delays involved in repeatedly de-/re-tunnelling, and also robustness issues (see later).

Much prior work has now been merged into two Internet Drafts, one for mobile IPv4 and one for v6, which are now discussed.

6.8.2 Regional Registration for Mobile IPv4

Regional registration introduces a Gateway Foreign Agent and optionally Regional Foreign Agents as level(s) of hierarchy below the GFA. The mobile can use either a co-located or foreign agent care-of address (i.e. as normal). This is called the 'local CoA' and is registered with the GFA (or RFA, if present), whereas the GFA's address is registered with the home agent as the mobile's CoA.

The foreign agent includes the 'I' flag in its advertisement, to indicate that regional registration is operating in the access network. The advert announces the GFA's address as well as the FA's address (or its NAI).

Two new message types are introduced: the regional registration request and regional registration reply. These are just like the normal MIPv4 registration request and reply, but are used for registration with the GFA.

6.8.3 Home Registration

When the mobile changes GFA or arrives in a new access network, it performs a 'Home Registration'. This involves sending a MIPv4 Registration Request to the GFA [14], with the care-of address field equal to the GFA address, and with the mobile's CoA included in the Hierarchical Foreign Agent extension. The GFA updates its visitor list and then sends the registration request on to the home agent.

6.8.4 Regional Registration Request

When the mobile moves to a new FA but is still on the same GFA, it performs a 'Regional Registration'. This involves sending a Regional Registration Request to the GFA, informing the GFA of its new 'local CoA'; the GFA does not inform the home agent. If RFAs are being used, there is one extra complication, which is that it is possible for the tunnels to become incorrectly directed if a mobile moves back to a previous FA. The solution is to explicitly deregister the old entries by sending a binding update with a zero lifetime to the mobile's previous CoA.

Thus, after the initial home registration with the home agent, subsequent mobile registrations can be localised within the access network.

6.8.5 Hierarchical Mobile IPv6

This is very similar to regional registration, and most of the differences are terminological - for example, the local mobility agent is called the Mobility Anchor Point (MAP). There are in fact two modes.

In 'basic mode' hierarchical mobile IP, a mobile obtains its CoA (called a Regional CoA, RCoA), through standard stateless address autoconfiguration (i.e. it consists of the MAP's subnet prefix plus the mobile's interface identifier). This is a globally routable address that the home agent (and correspondents, after route optimisation) routes to; in other words, the RCoA routes packets as far as the mobile's MAP. In order to route packets the rest of the way, each time the mobile moves, it sends the MAP a binding update with its new 'on-link' care-of address, LCoA. The message is just a regular mobile IPv6 binding update with an extension, the M flag, to distinguish it from a regular home registration or route optimisation message.

Mobiles need to discover nearby MAP(s). One method is for the MAP to send out a MAP router advertisement (which is a regular IPv6 router advertisement, with an extension containing the MAP's global address). This propagates through the access network until it reaches the access routers, which transmit it over the air. Hence, a mobile can listen to the advertisements and work out when it has moved into a new MAP's area; the mobile can then obtain a new RCoA and send regular mobile IPv6 binding update(s) to its home agent and correspondent(s). Further extensions can indicate the MAP's 'preference' (so an overloaded MAP could lower its preference rating, for instance), and the number of hops to the MAP. The last feature might let a mobile choose a distant MAP if it is moving extremely fast, to reduce the frequency of inter-MAP updates, and a nearby MAP if it is communicating with a local correspondent as suggested in [Figure 6.5](#).

The other mode of hierarchical mobile IP is called 'extended mode'. This deals with the scenario where there are mobile routers, i.e. a mobile that has further mobile hosts attached to it - for example a Personal Area Network. The idea is that the mobile router acts as a MAP; the mobile hosts therefore use the mobile router's RCoA as their CoA, called the alternate CoA. This is like the Foreign Agent variant of mobile IPv4.

6.8.6 'Fast and Smooth' Mobile IP-based Schemes

Mobile IP, as described above, can suffer from a break in communications during a handover. This section will outline ways specifically targeted at making the handover smoother (i.e. packets are not lost) and faster (i.e. packets reach the mobile with a smaller delay). Unlike the local mobility techniques above, these are not concerned with reducing the signaling load.

The basic approaches that can be employed are described below.

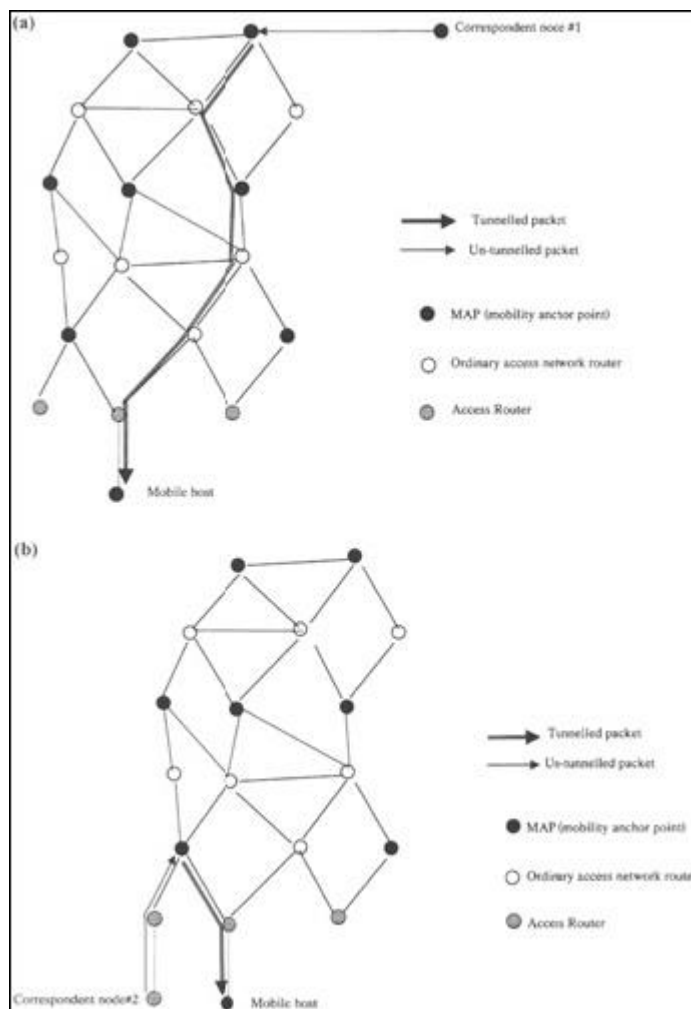


Figure 6.5: Hierarchical mobile IPv6. A mobile node selects a different MAP for correspondent nodes #1 and #2, as shown in (a) and (b) respectively.

6.8.7 Two CoAs

This is actually a feature of Mobile IP, and applies if the mobile is capable of listening on two links at once (i.e. make before break). Providing a mobile is allowed to hold on to its old CoA for a short period of time after the handover, it can accept packets arriving at the old or new link. This means that when a mobile hands over, packets that are sent before the binding update reaches the home agent (i.e. to the old CoA) will still reach the mobile and be accepted by it. The new 'fast mobile IP' schemes (sort of) extend this idea to break before make handovers. A mobile can configure its new CoA whilst still attached to the

old access router - this speeds up the registration process when the mobile moves on to the new access router.

6.8.8 Simultaneous Bindings and Packet Bi-casting

This is an optional feature in Mobile IPv4. A mobile sets the 'S' flag in its registration request, and the home agent interprets this as a request to retain the previous mobility binding(s), as well as adding the new binding (hence 'S' for simultaneous bindings). Subsequent packets from a correspondent can then be duplicated ('bi-casted') by the home agent, with a copy sent to each CoA.

This idea has been extended by performing the duplication locally (e.g. at the foreign agent). Additionally, it has been proposed to buffer packets locally during a handover.

6.8.9 Temporary Tunnel

This is another optional feature of Mobile IP that 'fast mobile IP' schemes propose extending. The basic idea is to establish a temporary tunnel from the previous CoA to the new CoA. Hence, packets coming from correspondent nodes that have not yet been told the new CoA (or indeed packets in flight from the home agent) will be forwarded on to the mobile. In Mobile IPv6, the method is that, whilst at its old point of attachment, the mobile discovers (from router advertisements) a 'local' router that has the capability to act as a home agent (typically just the old access router). Now, when the mobile connects to the new link and receives its new CoA, it sends a binding update to this router with a special field set (the home registration bit), which asks the router to act as a temporary home agent - it can then intercept packets addressed to the old CoA and tunnel them on to the mobile at its new CoA. The same idea is seen in Mobile IPv4's Route Optimization extension. The mobile adds the Previous Foreign Agent Notification extension to its binding update, which causes the new FA to send a binding update to the previous FA, which sets up a tunnel between the previous FA and the new CoA. (A FA indicates that it can support such forwarding by setting the 'S' flag in its agent advertisement. Here, 'S' is for smooth handovers.).

We see below that the method has been extended for the case where the tunnel can be set up in advance of the actual handover.

'Fast mobile IP' schemes are under intensive development in the IETF's mobile-IP working group. Many protocols have been proposed, but work is now converging into one protocol for IPv4 and one for IPv6. A few details are now outlined, although some are liable to have changed by now.

6.8.10 Fast Handovers for Mobile IPv6

The main idea of this protocol is that it is often known what the next Access Router is likely to be before the mobile actually hands over to it. This 'hint' could, for example, come from power or signal-to-noise ratio measurements, or from knowledge of a mobile's likely movements (e.g. if it is on a train). Hence, some proactive action can be taken in advance of the actual handover - if desired, the handover can be initiated before the MN has connectivity with the new AR. Overall, this should mean that from the point of view of ongoing communications between the mobile and its correspondents, the handover is apparently smoother and faster. This type of approach is familiar from current cellular

networks, where the mobile reports on the signal strength from nearby base stations, thus allowing the network to plan for handovers.

The basic ideas are to:

- Enable the mobile node to configure a new CoA before it moves on to the new AR, so that it can use the new CoA as soon as it connects with the new AR. This eliminates the delay seen in mobile IP from the registration process, which can only begin after the Layer 2 handover to the new AR is complete. There is an implicit assumption that the mobile is only capable of connecting with one AR at a time, i.e. break before make, otherwise the mobile IP feature above (two CoAs) can be used.
- Ensure that no packets are lost during the handover by establishing a temporary tunnel from the old to the new AR. The technique is basically an extension of the MIP feature above (point 3) to the case of a planned handover.

Figure 6.6 shows the basic messages involved.

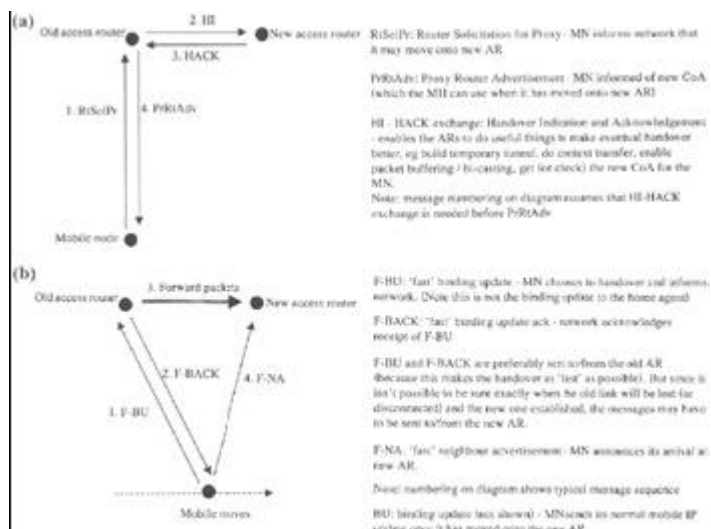


Figure 6.6: Fast mobile IPv6 handover. (a) Handover preparation phase of mobile-controlled scenario (b) Handover execution phase of mobile-controlled scenario.

6.8.11 Fast Handovers for Mobile IPv4

Fast handovers (or 'low latency handoffs' in their terminology) have also been considered for mobile IPv4. At present, there are several differences from fast mobile IPv6, and in fact, fast mobile IPv4 currently includes two different techniques called pre- and post-registration. Fast mobile IPv4 and v6 might be expected to converge on the same basic approach.

The 'pre-registration' method has the same idea as fast mobile IPv6 above, in that a proxy router advertisement from the old foreign agent is used to inform the mobile node about the prospective new foreign agent. There are several slight differences, which mainly stem from using a normal registration request/reply (i.e. there are not special 'fast' messages).

The 'post-registration' method is slightly different. It is more like normal mobile IP, in that no attempt is made to register the mobile node with the new FA until after the mobile has a Layer 2 link established with it. Instead, some sort of Layer 2 trigger causes the network to set up a 'bi-directional edge tunnel' (BET) between the old and new FAs. The old FA bicast packets to the new FA down the tunnel, so that when the mobile node makes a Layer 2 connection with the new FA, it immediately obtains its downlink packets. It can also send packets immediately - still using its old care-of address as the source address - because the new FA tunnels them to the old FA, where the packets are de-capsulated and forwarded; if the tunnel were not in place, the new FA might filter the packets because the source address was suspicious. Meanwhile, the mobile can, at its leisure, use standard mobile IP to register the new CoA; subsequently, it will of course need to tell the old FA to stop bicast and to tear down the tunnel.

6.9 Per-host Forwarding Protocols

6.9.1 Outline of their Operation

These protocols use a new specialised scheme to install per-host forwarding. The general idea is that information is stored in various routers spread through the access network. A downstream packet enters the access network (AN) at the gateway. The gateway looks up which of its output ports is the best to use, for the particular mobile in question (hence the term 'per-host forwarding'). It then forwards the packet on the selected port towards the 'next hop router'. At that router, the process is repeated, i.e. it in turn selects the best output port for this mobile. Eventually, the packet will reach the access router (AR) to which the mobile is attached. Thus, we see that:

- Information about the mobile's location is distributed throughout the access network.
- Packets are forwarded to the mobile without tunnelling or address translation.

Thus, a mobile keeps its address whilst it's within the access network-this is a major contrast to the tunnel-based schemes (covered earlier), i.e. the mobile does not have to obtain a new care-of address each time it moves on to a new access router.

The major job of a per-host protocol is therefore to:

Distribute (i.e. initialise) the forwarding information in the various routers.

Maintain the forwarding information.

- Update the forwarding information as the mobiles move. An important concept is the 'cross-over router', that is the router where the paths to the old and new access routers diverge.

So, when a mobile hands over, the cross-over router (at the very least) must change its perhost forwarding entry. This is achieved by the mobile sending a route update message when it moves, which installs the new entry(s) as required.

In general, per-host schemes hope that by confining signaling to a local region, i.e. near the access routers and the cross-over router, the signaling load will be reduced compared with basic mobile IP, and also that the handover will be much smoother. How effectively

a particular protocol achieves these aims will depend on the network topology as well as the details of the protocol. The protocols also have various extra techniques to try and achieve smooth handovers. In general, upstream packets simply go on the default route, towards the correspondent.

Three different sorts of per-host forwarding protocols for IP micromobility are discussed below.

6.9.2 New Protocols - Cellular IP and Hawaii

Initialisation of the forwarding information is done using reverse path forwarding: when a mobile turns on or enters the access network it sends a packet on the default route (i.e. shortest path) to the gateway; each router caches an entry mapping the mobile's identifier (its home address) to the neighbour from which the packet arrived at the node. Thus, downstream packets can be delivered to the mobile simply by following the series of cached mappings associated with that mobile, i.e. reversing the path.

Forwarding entries are soft state, i.e. they need to be periodically refreshed or, after a while, they will time out and be deleted.

Handover detection is simple and similar to mobile IP. It is done at Layer 3, i.e. it relies on the mobile 'hearing' the advertisement from a new access router. It then establishes a new connection and sends an update message. The 'trick' seen in fast mobile IP of using Layer 2 information to trigger action before the actual handover is not done - though presumably it could be added.

Some details of Cellular IP and Hawaii are now discussed. The operation of Cellular IP is Outlined in [Figure 6.7](#) and of Hawaii in [Figure 6.8](#).

6.9.3 Cellular IPv4

In addition to the general outline described above, specific features of cellular IPv4 are:

- The mobile is identified by its home address - As far as mobile IP is concerned, the gateway acts as the mobile's foreign agent and so the gateway's address is used as its CoA.
- Mobile to mobile calls are routed via the gateway, even if a more direct route exists.
- The route update packet travels all the way to the gateway, installing entries up to the cross-over router and refreshing entries above this (i.e. nearer to the gateway). This ensures that downstream packets follow the shortest path route to the mobile. Old entries, i.e. between the cross-over node and the old AR are simply left to time out.

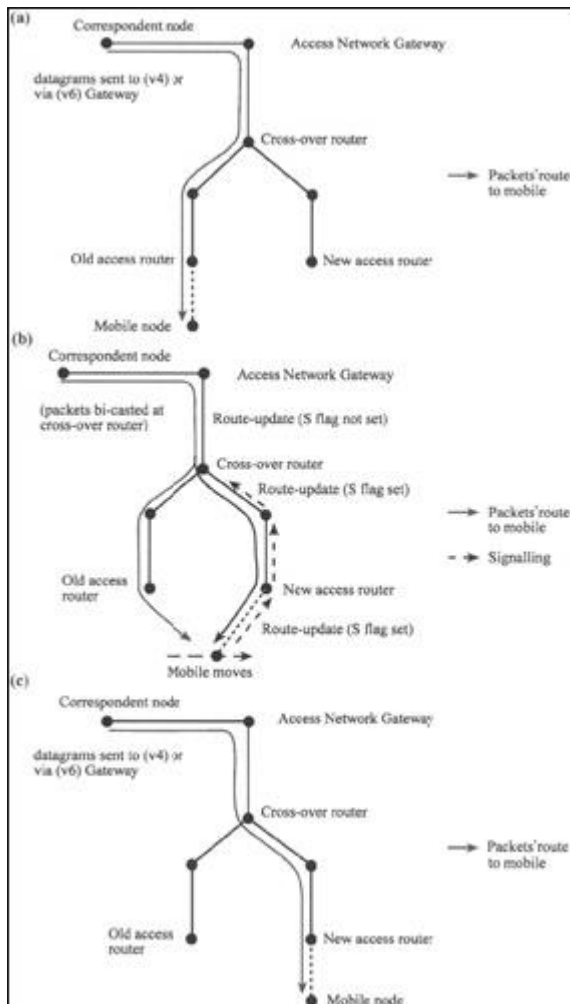


Figure 6.7: Cellular IP. (a) Before handover (b) During handover (bracketed comments apply to semi-soft handover case: once route update message reaches cross-over router packets are bi-casted to both old and new access routers.) (c) After handover.

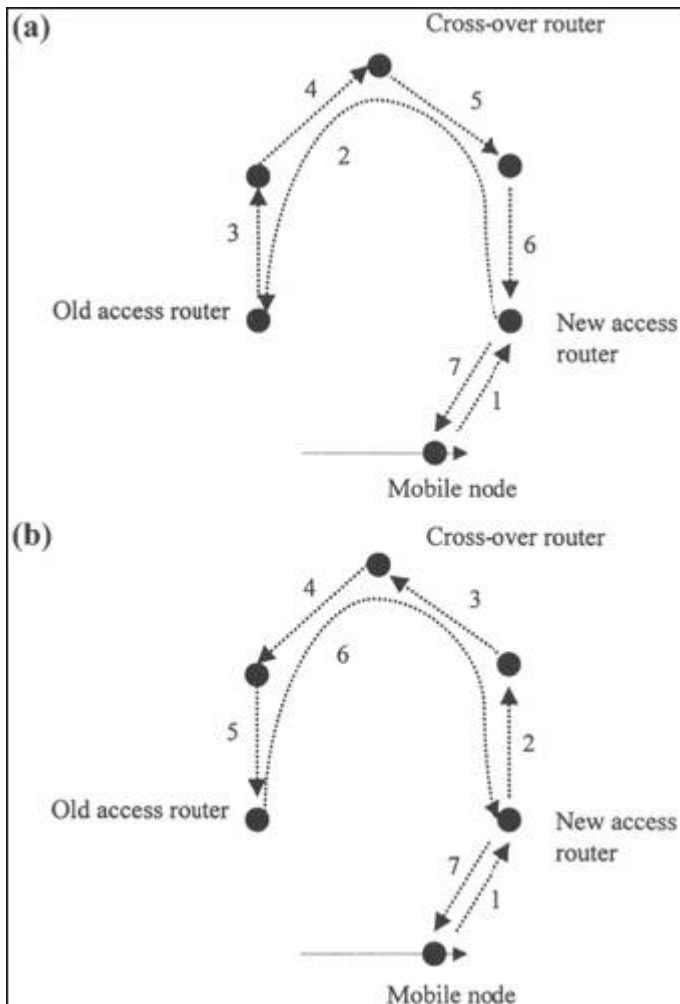


Figure 6.8: Hawaii handover messaging. (a) Forwarding scheme (b) Non-forwarding scheme.

- Cellular IPv4 can use ordinary data packets as implicit refresh messages. (An earlier idea was that route creation and updating were also done implicitly. However, this means that a router must 'snoop' all data packets, which is generally considered undesirable for security and performance reasons.)
- Cellular IPv4 actually refers to 'cellular IP nodes' rather than routers. The idea is to make a device that is a slimmed down router and so is cheaper. This no longer seems to be viewed as a likely benefit, so the terminology of 'routers' here remains.
- There is an alternative type of handover, called 'semi-soft', that aims to make the handover more seamless and is applicable when the mobile can listen to transmi from the old AR at the same time as sending to the new AR. The method is basically the same as simultaneous bindings in mobile IP; the mobile sets a flag ('S') in the route update packet, which the cross-over node interprets as an instruction to forward downstream packets to both the old and new ARs.

6.9.4 Cellular IPv6

Cellular IPv6 updates cellular IPv4 with IPv6 capabilities and adds a couple of minor changes of which the most notable are:

The mobile is identified by its (co-located) care-of address, which it keeps whilst it is in the AN, so there is no need for the gateway to act as a foreign agent (contrast Cellular IPv4). The CoA is obtained through IPv6 stateless autoconfiguration (i.e. CoA is the gateway's IPv6 subnet prefix plus the mobile's interface identifier).

- Another alternative handover has been added, called 'indirect semi-soft'. This assumes that a MH cannot listen to the current AR whilst sending a route update packet to the new AR (as required by semi-soft handover). Instead, when a MH decides to handover, it sends an update packet to the current AR. This packet's destination address is the new AR, and it has the 'T' flag set. When the new AR receives this packet, it interprets it as an instruction to create a normal semi-soft update packet.

Hawaii

Again, the protocol follows the principles given in the Outline section above. There is no Hawaii for IPv6 as yet. The specific features of Hawaii are:

- The mobile obtains a co-located care-of address, when it is not in its home domain, which it keeps whilst it is in the same Hawaii domain.
- Mobile to mobile calls are routed on the most direct route available, i.e. not necessarily via the gateway.
- The mobile sends ordinary mobile IP registration messages. At the access router, these trigger Hawaii messaging inside the domain. The aim is that the operation of Hawaii is hidden from the mobile.
- Route updates only travel as far as the cross-over router.
- Explicit refresh messages are generated hop by hop, thus allowing for their aggregation.
- There are two different approaches for how the network reacts to a handover. The 'forwarding scheme' is appropriate when the mobile can be connected to only one AR at a time. It results in downstream data packets being first forwarded from the old AR to the new AR before they are diverted at the cross-over router. However, the 'nonforwarding scheme' is appropriate when the mobile can be connected to both ARs simultaneously. Downstream data packets are diverted at the cross-over router as soon as the path update message reaches it, and so there is no forwarding of packets from the old AR.
- Note that, because route updates are directed from the new AR towards the old AR (rather than to the ANG as in cellular IP), this means that after several handovers, the path taken by the downstream packets may not be the most direct available. An example is shown in [Figure 6.9](#).




Figure 6.9: After several handovers, route may be less direct in some per-host forwarding schemes.

MANET-based Protocol - MER-TORA

Currently, there is only one proposal in this category: MER-TORA. Its operation is outlined in Figure 6.10. MER-TORA builds on the TORA ad hoc routing protocol.




Figure 6.10: MER-TORA. The height shown is with respect to destination, i.e. mobile host. Only two elements of the height are shown: the reference level (which is

decremented to indicate that the mobile has moved); and the delta (showing the hop count from the reference level). Heights are ordered first by reference level, then by delta. (a) Before handover (b) During handover (c) After handover.

In TORA, each host and router has a 'height' associated with it. A packet is routed downhill from a source to its destination. The TORA protocol assigns all nodes an appropriate height and then reacts appropriately to any changes in routing topology (e.g. a link failure) to ensure that there is still a downhill route to the destination. Note that a 'height' is assigned with respect to a particular destination, i.e. a separate version of the protocol is run for each destination. The 'height' has five elements to it, but in the case of a static network, a node's 'height' is essentially its hop count to the destination. A node can have several downhill links to the destination, which means that TORA copes elegantly with meshed (non-hierarchical) networks; it is the job of some other protocol to decide between alternative routes (maybe based on QoS). If a link breaks, and this was the last downhill link to the destination, the TORA protocol automatically reacts by trying to discover a new route; the height then becomes more complex than a simple hop count. This route discovery is designed to be as local as possible, so it is claimed that routing will be restored very rapidly.

When applying TORA in an Access Network, a few changes are suggested. First, to reflect that the AN is much more stable than a MANET, TORA is run proactively. This involves the occasional propagation of an optimisation packet through the network to re-initialise (reoptimise) the routing by flushing out the effect of any node/link failures. Second, an efficient address allocation mechanism is suggested. The idea is that each AR owns a block of IP addresses and that the TORA algorithm is run on this address prefix. When a MH turns on and attaches to an AR, it is assigned one of its addresses; once this has been registered (e.g. in a SIP location database or at a MIP home agent), packets can be routed to it using TORA. Thus, this prefix-based routing allows the ARs and any static terminals to be reached. Third, something must be done when a MH moves, since the prefix-based routing to this MH will no longer work. One solution would be to run TORA for the MH, but this would install a hostspecific route for the MH throughout the AN. Instead, we would like to use prefix-based routing through most of the AN and just add some host-specific routing near the 'edge' of the AN. An extension to TORA has been devised to achieve this. The mechanism is called MERTORA, and the idea is to send a UNICAST_UPDATE packet from the new AR to the old AR (along the prefix-based TORA route), which installs host-specific entries as it travels (the entry is just the node's new height with respect to the MH).

Next, when the mobile 'switches off', the IP address is returned to the allocating-AR, and the host-specific routes are deleted (essentially, this is treated as a handover to the original AR).

Finally, it is also suggested that in a planned handover, the MH should inform the old AR to which AR it is about to handover. The old AR can then build a temporary tunnel to the new AR in advance of the actual handover. This enables the re-direction of packets that would otherwise be lost in flight whilst the new host-specific route is being installed. The idea is the same as that for fast mobile IPv6. Similarly, a virtual path between the two ARs may be used to swap messages (e.g. warning of the impending handover,

exchanging information on available resources or authentication details, etc. See also Context Transfer later).

6.9.5 Multicast-based Schemes

Multicast protocols are designed to support point-to-multipoint connections, for instance to distribute Internet radio or TV to interested people. The basic principle in using multicast for mobility is therefore to assign a multicast address to a mobile and, when it moves, to add the new access router as a new leaf on to the multicast tree and remove the old AR (either by explicitly pruning it off or by waiting until its soft state multicast entry times out). If it is known that a handover is imminent, this can be done in advance, thus making the handover seamless.

One idea is that the multicast address assigned to the mobile should be public, i.e. globally routable, but this would require large-scale management of multicast addresses across the public Internet, which is unfeasible.

A more plausible idea is to keep the multicast addresses private to the access network and use the gateway's address as a care-of address - so, for example, it would be registered at the mobile's home agent, effectively as a foreign agent CoA. This enables the use of multicast to be hidden from the wider Internet, but does entail the gateway acting as a foreign agent, i.e. decapsulating downstream packets to discover the mobile's home address, looking up the appropriate multicast address, encapsulating the original packet inside a multicast packet, and then sending it into the multicast tree. The AR then decapsulates the packet and delivers it to the mobile.

There are two categories of multicast protocols - sparse mode and dense mode - and mobility protocols have been proposed, based on each.

However, such proposals have largely met with hostility. Several reasons are suggested:

- Functional overload - It is not what multicast was designed for. In particular, multicast protocols are not particularly good at dealing with quick changes of multicast group membership, which is exactly what a fast handover requires.
- Although a multicast-based protocol easily allows packets to be duplicated, so that they are ready and waiting at the new AR, as has already been seen, this functionality can be easily incorporated into any of the other IP mobility approaches.
- Although multicast-based schemes fall under the per-host forwarding banner, they also use tunnelling. Thus, they might be expected to suffer from the disadvantages and difficulties of both.

Multicast-based schemes are ignored in the comparison section that follows next.

Using SIP for micromobility would raise similar problems. Note also that the operator may not want to be driven by mobility considerations when positioning SIP servers and SIP location databases in the network.

As will be seen later, this does not apply to all per-host forwarding schemes.

A flag is a particular bit in the header that the protocol defines to have a particular meaning.

It is almost possible simply to use the normal MIPv4 registration request and reply, but unfortunately, there are a couple of detailed optional cases where it would not be possible to distinguish between whether a regional or normal registration was intended. If the MH has a CCOA, it can send the registration directly to the GFA. However, if it has a FA-COA, the registration must be relayed via the FA.

An option is to set the CoA field to zero, in which case, the mobile is assigned a GFA.

Trying to run basic mode in this mobile router scenario is fraught with difficulties. For example, if the mobile hosts obtain their own RCoA from the mobile router (i.e. their MAP is at the mobile router), then every time the mobile router moves, it has to obtain a new network prefix, and the mobile hosts have to obtain a new CoA. By contrast, if the mobile hosts obtain their RCoAs from a MAP further in the network, then after the mobile router has moved, their RCoAs will no longer be globally routable.

Chapter 7

Comparison of Micromobility Protocols

7.1 Comparison of Micromobility Protocols

Much comparison between the various micromobility protocols is contentious after all, the proponents of every protocol believe that theirs is the best because it compares favourably with the others. This section tries to be fairly neutral, discussing the key issues architecture, scalability, reliability, and philosophy (or implicit assumptions). This means that some people's pet subjects are liable to be ignored.

The first thing to point out is that the protocols are much more similar than is normally admitted. Indeed it would seem not unreasonable to say that the protocol chosen is largely a matter of taste; of course, there are differences, but in many circumstances, these will not mount to much. Indeed, their similarity is being strengthened by a process of merger and acquisition. This is the nature of the IETF standardisation process; various protocols are suggested, and similar protocols are gradually being merged, whilst good ideas are gradually introduced from other proposals. This has led to some architectural convergence, i.e. some agreement on the architectural principles that are a good thing, as described below.

7.2 Operation

At the bird's eye view, as [Figure 7.1](#) and [Table 7.2](#) show, all the protocols do the same thing in the same way but use different names for the messages and nodes. The key idea is that path updates are localised - they only travel between the cross-over router and the old and new access routers (ARs). This minimises the signaling load and also ensures that the path update process is reasonably rapid.

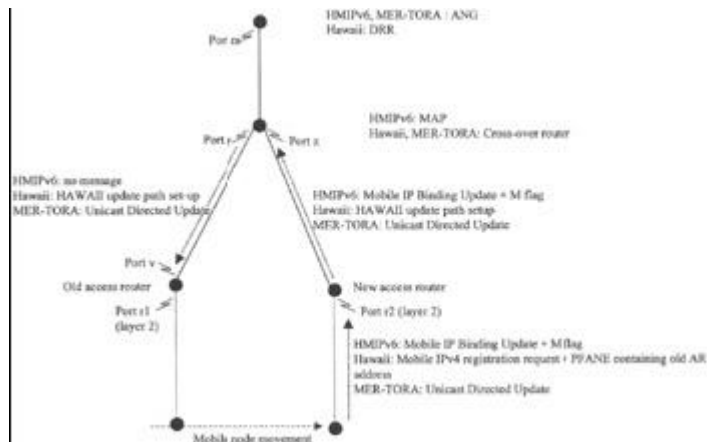


Figure 7.1: Comparison of selected micromobility protocols, showing messaging during

handover. MER-TORA: default routing entries point towards address allocating router – have assumed this is old AR. Also assumed mobile IP for global mobility. HMIPv6: default routing entries point towards MAP (Mobility Anchor Point). Assumed basic mode, so Regional CoA (RCoA) identifies MAP. Hawaii: default routing entries point

towards Domain Root Router (DRR), which is Hawaii's terminology for Access Network Gateway (ANG). Assumed nonforwarding variant of Hawaii.

7.3 Architecture

There are several 'architectural' choices that an IP micromobility protocol must make. Here, we look at four:

- Is it better to use link layer or IP layer information to detect when a mobile has moved on to another access router?
- Should handover be mobile- or network-controlled?
- Should handover be made smoother using packet buffering, or bi-casting, or both?
- How many separate protocols should there be to manage mobility?

This section gives the best estimate of the way things are going - there is considerable emerging consensus, although some of the points below are somewhat speculative.

7.3.1 Handover can be Initiated Using Layer 2 'Triggers'

The issue here is how to perform 'movement detection', i.e. decide that the mobile is moving on to a new Access Router (AR), and so a handover should be initiated. Basic mobile IP does this using messages at the IP layer. This has the advantage of giving a clean separation between Layer 2 and 3, so that mobile IP can operate over any link layer. However, it does increase the handover latency because the Binding Update can only be sent once the Layer 2 handover is complete, and the advertisement from the new AR has been received.

Fast MIPv6 and MER-TORA have introduced the idea of using information available at the link layer in order to trigger the Layer 3 messaging in advance of the actual handover, and thus reduce its latency as perceived by an IP user. This information could originate from the mobile (mobile-initiated handover) and be based on measurements of signal-to-noise-ratio from nearby base stations or whatever is best for the particular radio technology. The information might also originate from the network (network-initiated handover) (for example) when the network realises that a cell is becoming overloaded, and so it should do some load balancing by initiating a handover.

An extension of this idea is to standardise the interaction between the link and IP layers, by defining a generic interface between them. Clearly, the interface would include some sort of identifier for the potential new AR and perhaps also an indication of the urgency of the handover. Defining this interface in a sufficiently generic way, including being able to cope with future wireless technologies, is non-trivial. Such a generic interface would also help with inter-technology handovers.

7.3.2 Handover is Mobile-controlled

In traditional cellular systems, the handover is network-controlled. This is because the operators like complete control over their network, chiefly so that they can optimise the use of scarce radio resources. By contrast, mobile-control is assumed by mobile IP and all the (current) micromobility protocols, meaning that the mobile is in ultimate control of the handover, in terms of deciding exactly when it will handover. The main reason is that the mobile is best placed to understand the user's needs in terms of all their varied applications, the competing requirements from operating system activity, and their

multifarious personal preferences (e.g. to choose a higher-quality, but more expensive, link). This is essentially an instantiation of the end-to-end principle of IP design. Traditional cellular systems do not need to give the user ultimate control, because they only provide a single service, voice, that is well understood, so it is easy for the network to give the user what they want.

Although the handover is mobile-controlled, the network can initiate, assist (see context transfer later), or constrain (e.g. reject) it. Further, there may be some scenarios when there is a requirement to approximate a network-controlled handover (e.g. the mobile is too simple to decide about a handover). This could be achieved by a network-initiated handover with an 'urgent' flag set - if the mobile ignored this, it would find that the network had cut its connection. Some people believe that this will not work, and that network-controlled handover is required in order to achieve effective real-time quality of service and radio resource management.

7.3.3 Packet Bi-casting and Buffering

Packet bi-casting and buffering are features that the access network may need to support. However, several things are not clear:

- Where is the best place to perform bi-casting (i.e. packet duplication) - at the old access router or at the cross-over router?
- Where is the best place to perform buffering (i.e. packet store-and-forward) - at the old access router or at the new access router?
- Should these two features be the default or strictly optional?

The answers may depend on the type of application and the general scenario. For example, if the mobile is rapidly ping-ponging between two ARs, bi-casting looks attractive. Buffering may be preferable if bandwidth at the edge of the network is seriously constrained, or if it is critical that there is no packet re-ordering. If it is important to minimize the gap during a handover (e.g. for a critical real time application), it may be better to buffer at the new AR than the old AR.

These sorts of questions and answers are currently under active disagreement at the IETF. One solution could be to add flags to the handover protocols that allow mobiles to request either buffering or bi-casting, and for the network to choose where (and whether) to support these features. This would give the flexibility to the users and network operators to make the decisions based on their particular scenario

7.3.4 Separate 'Handover' from 'Path Updates'

In MER-TORA and fast MIPv6, the signaling during a handover only involves the mobile node and the old and new ARs; it is only once the mobile has finally and definitely moved on to the new AR that an update message is sent into the network towards the cross-over router. Hence, the approach is sometimes called 'edge mobility'. By contrast, cellular IP, for example, does not make this distinction, and a handover immediately triggers a message (route update) into the network.

The 'separation' approach seems particularly appropriate for planned handovers, where it is uncertain whether the mobile will actually move on to the prospective AR. Confining

signaling to the 'edge' of the network makes it easier to fall back - there is no state to find and flush out somewhere within (perhaps deep within) the network. It also seems easier in practice to devise signaling that can deal more easily with failure cases, like loss or misordering of messages.

Note that the separation of handover and path update signaling suggests a way to combine (for example) fast Mobile IPv6 and hierarchical mobile IPv6: the former is used for the handover signaling, and the latter for path updates. This separation would mean that hierarchical mobile IPv6 would not have to ensure a seamless handover - that is a requirement that handover management deals with - and conversely, 'Fast Mobile IPv6' would not have to deal (much) with scalability.

A further possibility is to standardise the handover signaling, but not the path updates. This would give a universal standard for signaling over the air and thus would allow a mobile to work on any network, providing, of course, that it had the appropriate wireless link technology. But it would also allow an operator to choose a path update protocol appropriate to their particular circumstances.

7.3.5 Separate Paging Protocol

'A similar separation' question is whether paging (see later) should be a separate protocol. The IETF's Seamoby working group recommends that it should be, whereas several existing protocols, like cellular IP and Hawaii, incorporate paging and handover within one protocol. Advantages of protocol separation might be that it will lead to a more fundamental consideration of the best way to do paging, and that it may allow an operator to deploy the Seamoby paging solution, but their own mobility protocol.

7.3.6 Separate Macro and Micro Mobility

Another question is whether macro- and micromobility management should be kept separate or integrated together. The advantage of separating them is that the access network operator does not have to worry about what macromobility protocol is used (it could be MIP or SIP or indeed nothing), which means that the two protocols can be deployed and evolve independently. It perhaps also allows cleaner security and authentication. On the downside, there may be some extra signaling over the air (because there is no possibility of combining the messages), and also, it is much harder to deploy foreign agents (which may be relevant if IPv4 addresses are at a premium).

Most micromobility protocols to date have assumed mobile IP for macromobility: cellular IP, Hawaii, hierarchical mobile IP, and regional registration all do so. A couple of protocols that do not are BCMP and IDMP (see references) - they are otherwise rather similar to hierarchical mobile IP, although BCMP goes on to point out that any tunnelling protocol can be used, not just mobile IP (Figure 7.2).

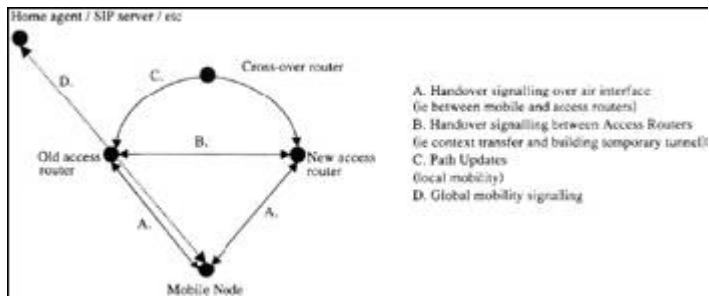


Figure 7.2: Possible separation of handover-related signaling into different protocols.

7.4 Scalability

"Scaling is the ultimate problem many ideas quite workable in the small fail this crucial test". Such is the IETF's clarion call to protocol designers.

For IP mobility, although, in some deployment scenarios, scalability will not be a concern (e.g. in the home or a small office), we would like our solution to scale to big networks, ideally up to global in scale.

7.4.1 Tunnel-based Schemes

Mobile IP is designed to scale - as the number of mobile nodes grows, more home agents can be added. Indeed, one idea is that a user's home PC (on an ADSL link, so 'always on') acts as the home agent as the user roams elsewhere with their laptop or personal communicator. Apart from this, all that is required is that an Access Router records the mobiles that are currently attached to it. However, although mobile IP scales in terms of information storage, it does not with regard to its messaging - since a handover could lead to long-distance signaling back to its home agent. This has led to the development of Regional Registration for mobile IPv4 and hierarchical mobile IP for v6; these reduce signaling by introducing a local mobility agent that essentially acts as a local proxy home agent. Movement within the local area (micromobility) only needs to be signalled to the nearby mobility agent. Scalability can be further assisted by introducing more mobility agents as the number of ARs and/or mobiles increases. However, movement between the local areas (macromobility) involves an update back to the home agent, so at some point, the introduction of further mobility agents will do more harm than good (at least in terms of signaling scalability). A further level of hierarchy could then be introduced: 'regional areas' that contain several 'local areas' with corresponding regional mobility agents. The idea could be extended to an arbitrary number of levels of the hierarchy. Note that a network operator deploying any of these schemes must decide where to place the mobility agents and how many to use - not a trivial optimisation problem - and what to do as the network grows.

7.4.2 Per-host Mobility Schemes

A criticism often raised with per-host mobility schemes is that 'they do not scale'. Essentially, the claim is that they tend to distribute information about the mobile's location amongst many routers. To put this another way, it means that a router will have to store information about the location of many mobiles. The problem is likely to be most acute for the gateway: indeed, in cellular IP and Hawaii, the Gateway must have an entry

for every mobile within its access network. Clearly, this will eventually limit scalability, because as the network and number of mobiles grows, the forwarding table will grow, and eventually it will be too large to retrieve the information sufficiently quickly. However, it should be possible to cope with thousands of mobiles, so the solution would probably scale sufficiently for a campus network, say.

MER-TORA introduces an interesting idea to alleviate this problem. This is the use of prefixbased routing as much as possible, plus a per-host 'routing tail'. A mobile is assigned an IP address (a care-of address), based on its initial location, so that standard IP routing can be used to send packets there. When the mobile moves, information is added to just the few routers necessary in order to add a mobile-specific route on to the original route. For example, the first time that a mobile moves, this is probably just the two ARs and the cross-over router. Thus, whereas cellular IP and Hawaii use the IP address merely as an identifier, MER-TORA also uses it as a locator as far as it can.

WIP (Wireless Internet Protocol) has a similar idea of prefix-based routing plus a per-host routing tail, but it builds on a standard intra-domain routing protocol like OSPF. Initially, OSPF is used to do prefix-based routing to the MH's 'radio port'. When it moves, an update is sent to a local set of routers (the HARG, Handoff Affected Router Group), using reliable multicast. The HARG is unique for any pair of radio ports and could be calculated in advance. WIP maintains a shortest path to the MH, whereas MER-TORA only updates routers on the path between the new and old ARs; this is just like the cellular IP vs.

One issue is that gradually, the amount of per-host routing builds up as the mobile moves around, and eventually, there will be no scalability benefit. The implication is that the mobile occasionally must obtain a fresh IP address so that once again only prefix-based routing is needed and so the old per-host state can be flushed out. This should be done not just when the mobile switches off, but more often. However, to avoid disruption to the user, this needs to be done between sessions. This is easier said than done - it may be tricky for a mobile to tell when it is in between sessions, and it will also raise issues for a mobile that wants to provide an 'always-on' service for other users and so wants to be reachable via a permanent IP address (e.g. it is acting as a server). Provided that the latter is relatively rare, it could be dealt with as a special case.

One issue is that gradually, the amount of per-host routing builds up as the mobile moves around, and eventually, there will be no scalability benefit. The implication is that the mobile occasionally must obtain a fresh IP address so that once again only prefix-based routing is needed and so the old per-host state can be flushed out. This should be done not just when the mobile switches off, but more often. However, to avoid disruption to the user, this needs to be done between sessions. This is easier said than done - it may be tricky for a mobile to tell when it is in between sessions, and it will also raise issues for a mobile that wants to provide an 'always-on' service for other users and so wants to be reachable via a permanent IP address (e.g. it is acting as a server). Provided that the latter is relatively rare, it could be dealt with as a special case.

7.4.3 Soft vs. Hard State

As well as the amount of state, the amount of signaling to maintain that state can also be considered. Within the per-host family, there is an interesting contrast between Hawaii, which aggregates (merges) messages as they travel up the tree, and MER-TORA, which does not bother with any refresh messages. MER-TORA can do this because it uses hard-state routing, i.e. the path update messages are trusted as accurate and complete. However, Hawaii uses soft-state routing, i.e. it is accepted that some messages may be lost, and it may not always be possible to 'prune' off old routes. Clearly, MER-TORA must make more effort to ensure that the original path update is correctly received by all necessary routers - for example, each message is acknowledged. (See [Reliability](#) section.) Probably, either protocol could be modified to be either hard or soft state, so it is not a fundamental distinction.

7.4.4 Address Allocation

Scalability is also connected to address allocation, since the mobile needs to be allocated a care-of address. This is a particular issue for IPv4 networks, where the number of addresses (especially publicly routable ones) is limited. Cellular IP uses foreign agent care-of addresses and so is attractive if address shortage is critical. The basic message is that there are no ideal solutions to the IPv4 address shortage - except IPv6. Most per-host forwarding micromobility protocols require address allocation to be co-ordinated across the access network, whereas in MER-TORA, each AR owns an IP address block from which it allocates addresses to MHs as required. This probably makes address allocation simpler (but will probably require more addresses).

7.4.5 Multiple Access Network Gateways

Scalability is also affected if a particular router acts as a focal point for traffic. The obvious example here is the gateway to an access network-all upstream and downstream traffic goes through it. This is potentially bad both because it may become overloaded, and also, in a network that covers a large area, it means that traffic must travel over long distance routes. The obvious solution is to have multiple gateways. MER-TORA intrinsically supports multiple gateways, but it is not clear that other per-host schemes can be modified to do so. (Hawaii and Cellular IP can have multiple gateways, but each mobile can only use and be contacted through a particular gateway.)

7.4.6 Paging

Paging is a technique that, amongst other things, can improve the scalability of the IP mobility protocol by reducing the number of messages sent. As in GSM, for instance, 'paging areas' are defined, which contain several ARs, and a 'dormant mode' for the mobile is defined, when it is not actively sending or receiving packets. The idea is that a dormant mobile only sends an update message when it moves between paging areas (whereas, of course, an active mobile sends an update message when it moves between any two ARs), resulting in a significant reduction in location update signaling. One consequence is that when a correspondent wants to communicate with a dormant mobile, the exact location of the mobile within the paging area must first be found. This involves extra signaling, which clearly needs to be less than the reduction in location update

signaling. There is a design trade-off, and in general, the network operator will adjust the size and shape of their paging areas in order to minimise the total amount of signaling. As well as reducing the amount of signaling, paging can also reduce the amount of state stored in the network, particularly by per-host forwarding schemes. For example, this might be done by all dormant modes having their location stored in just a single specialist router, whereas, of course, active mobiles will have their location information distributed amongst several routers.

7.5 Reliability

The concerns here are to be resistant to router and link failures, and to the loss of signaling messages.

Loss of signaling is typically dealt with through confirmed messages, i.e. using acknowledgements. In fast mobile IPv6, for example, there are limits on how many times a message can be re-sent; if an Ack is still not received, the process falls back to standard MIP.

Perhaps the major concern, however, is to avoid as far as possible single point of failures. Clearly, the failure of an Access Router will unavoidably disconnect all mobiles attached to it, unless there is overlapping coverage from another AR. The network topology may also be such that the failure of some other router inevitably causes the network to partition, i.e. the mobile to become disconnected. Again, there is no solution, other than to suggest to the operator that they re-design their network. However, normally, the network design will mean that there is sufficient redundancy such that an alternative route does exist. The challenge therefore is to find an alternative route as quickly as possible, preferably so that there is no noticeable break in ongoing sessions and so that no packets are lost.

7.5.1 Tunnel-based Mobility Schemes

There are two cases to consider.

First, there can be the failure of a mobility agent, i.e. a tunnel end point. This disconnects the downstream traffic to all mobiles using a tunnel terminating there. Now, this sometimes may not be too bad - take the case where a user's home agent is their home PC; if it fails, only the user is disconnected. The user would then have to use a different home agent, which would entail obtaining a new home address, informing others of it (e.g. correspondents and the SIP server), and having to re-start sessions. However, in HMIP and Regional registration, if the local mobility agent fails, many mobiles may be disconnected. The obvious solution is that when a mobile realises, it re-registers using a different mobility agent, but this will still take time, and perhaps more significantly, it is an open question as to how best a mobile can promptly realise that it has become disconnected.

Note that the general case where there is a hierarchy of mobility agents is worse from a reliability point of view, simply because there are more failure points. For example,

hierarchical mobile IP now recommends against forcing packets to be sent down through its hierarchy of MAPs, because it diminishes the robustness of IP routing between the highest MAP and the mobile.

An alternative idea is to concentrate on making the mobility agents extremely reliable (e.g. high-end machines and hot standby). However, engineering a component for high reliability is expensive.

The second case to consider is where some other router fails, or a link goes down, i.e. not a mobility agent. In other words, this is a router/link through which the tunnels pass. This will be dealt with by the standard recovery mechanism of the routing protocol (say OSPF) - so that there will still be the required tunnel between the two end points, but it will now follow a new route. It is likely that this will take some time to achieve, perhaps several seconds at least.

7.5.2 Per-host Schemes

Per-host schemes effectively distribute information about a mobile's location across several routers, so this might be expected to have an adverse effect on reliability. However, it is not that simple.

The first thing to consider is whether the protocol has any single point of failures. The obvious point in cellular IP and Hawaii is the (one and only) gateway. It is not clear as to how to use a kind of back-up gateway, since all upstream packets go on the default route through the gateway. By contrast, in MER-TORA, the access network can have any number of gateways, and a mobile is reachable through them all.

The second consideration is what happens when some other router or link fails. In Hawaii, the standard routing protocol (e.g. OSPF) uses its mechanism to detect a failure and update its default route entry. This can trigger a refresh of Hawaii's per-host entries to the new upstream router. MER-TORA also relies on the routing protocol to detect and route round failures - which, in this case, is TORA. Because TORA has been designed for unreliable ad hoc networks (MANETs), its mechanism should be fast and robust. It uses a highly localised flooding mechanism to find a route around the failure. It also has the advantage of intrinsically supporting meshed networks - so some routers and links can fail, and an alternative route will already be known.

A more fuzzy reliability issue is that MER-TORA is a new, moderately complex routing protocol - so an operator will need to gain confidence that it works properly in all circumstances and that they understand how to deploy, upgrade, and manage it.

7.6 Philosophy

This section discusses various issues that are essentially about how (whether) an IP mobility protocol affects other IP protocols, such as RSVP. (RSVP is a quality of service protocol). There are highly contentious technical arguments about various compatibilities (or non-compatibilities, depending on one's opinion), which are only hinted at below. The battles are basically between tunnelled and per-host protocols. War tends to be waged on the 'path update' part of the protocols and not the 'handover' part, although where a

temporary tunnel is used between access routers during a handover, the same arguments will apply, although in a weaker (temporary) fashion.

However, the debate is philosophical: what do users want an IP mobility protocol to be? This may seem a bizarre question - surely after all the description and discussion in the chapter, it is obvious that an IP mobility protocol is something that sends packets to mobile hosts despite their movement, in a way that is reasonably scalable and robust? However, the question is intended to address a deeper issue. The contrast is between:

- Tunnelled schemes - Which say that a mobility protocol is an add-on built on top of the standard routing protocol. Effectively, this hides the hosts' mobility from the routers, with mobility support confined to a few specialised nodes (i.e. the mobiles themselves and the mobility agents).
- Per-host schemes - Which say that a mobility protocol is just a routing protocol, albeit a new protocol. Effectively, this exposes the hosts' mobility to the routers.

The distinction leads to different knock-on consequences:

- Tunnelled schemes - Mobility will have an impact on other protocols, essentially on any protocol that keeps state in routers and assumes that the source or destination address remains constant, because mobility causes the mobile's care-of address to change.
- Per-host schemes - Mobility will not have such an impact on other protocols, because the mobility is treated as a topology change within the local area. A mobility protocol should look exactly like an ordinary routing protocol as far as everything else is concerned.

Someone who believes that, in the future, most devices will be mobile is more likely to believe that the knock-on effects will be very important, and so is more likely to be in favour of a per-host scheme. However, those in favour of tunnelling approaches may also believe in mobile-dominated future but think that the knock-on consequences will only have a 'second order' impact and are outweighed by other factors, such as the easy deployment of tunneling approaches (in particular, only a few nodes need upgrading). An IP 'purist' will also argue in favour of per-host schemes, because tunnelling is in tension with the end-to-end IP design principle: that is, Internet protocols are designed with end-to-end signaling in mind, so changes due to mobility are liable to induce unwanted end-to-end signaling - and this may also degrade the application.

Note, also, that from a practical standpoint, an operator opting for a per-host protocol will probably want to roll it out gradually, i.e. not to every router all at once (compare the deployment of IPv6 today). This would be achieved by deploying islands of routers with the new protocol in a sea of unenhanced routers, and connecting the islands through tunnels. So, even a fervent believer in per-host IP mobility will probably have to tackle tunnelling issues.

7.6.1 QoS Compatibility

This is probably the most famous issue. When QoS is established for a mobile's connection, the process installs information in routers in the network, telling them that when a packet to (or from) this address is routed, a particular set of QoS parameters

should be applied (e.g. treat it as high priority). The problem with tunnelling is that it hides the original header. The options are to decapsulate and examine the inner packet's header (but this is expensive in processing terms) or to have a choice of tunnels between the two tunnel end points with different priorities or allocated bandwidths (but this is complex to manage). Hierarchical MIPv6 schemes are (arguably) somewhat better, because the packets are (mostly) not tunnelled but instead use extension headers in the IP packet; in practice, this makes it easier for routers to examine QoS fields in the IP header.

There is also a question of how to integrate RSVP and mobility (see [Chapter 6](#)). RSVP sets up QoS based on the routing path to the destination address of the end node. With per-host schemes, the mobile keeps its CoA as it moves. This means that as soon as the mobility protocol has installed a route to the new AR, an RSVP PATH message can be triggered over the new part of the path. By contrast, in tunnelling protocols, the MH's CoA changes each time it moves, so the RSVP has to be re-set up over the entire length of the path. It should help to use reverse tunnelling, since then the CN is not exposed to the MN's change in careof address, and QoS only needs to be re-established between the MH and local mobility agent. However, overall per-host schemes are probably simpler to integrate with RSVP.

7.6.2 Web Caching

Caching is now extremely important in order to prevent servers overloading from too many 'hits' and to reduce network traffic. Tunnelling interacts with web caching because it can only be done outside a tunnel. Hence, for example, for hierarchical mobile IP with a single GFA, caching can only be done at the GFA. This considerably reduces the network operator's design flexibility in terms of how they position their caches.

Of course, this is only a bird's eye view. It also makes assumptions like there being a local mobility agent at the cross-over router.

Of course, if link layer triggers are not available, one can always fall back to standard IP messages to trigger handover.

IETF Guidelines for Conduct, Request for Comments: 3184, S. Harris, October 2001.

That is, the MH, rather than routing directly to the correspondent node, tunnels packets to the local mobility agent, which decapsulates them and forwards them on to the CN.

7.7 Other Aspects of Terminal Mobility

There is more to mobility management than the 'simple' problem of sending packets to the new AR and making the handover as seamless as possible:

- Context transfer - This concerns how the new AR learns about the 'context' associated with the mobile's communications session(s). An example of 'context' is the protocol state needed to carry out header compression over the air.
- Paging - This concerns mobiles that are not actively transmitting or receiving packets, and so can enter a dormant mode (sometimes called 'idle' or 'standby').

One reason is that a mobile in the dormant mode can conserve its battery power. The main challenge is to find and wake up ('page') a dormant mobile.

- Security - How to ensure that the mobile's communications are secure.
- Other points - For example, radio resource management, e.g. choosing between various alternatives for the new access router. These issues have had limited consideration by the IETF, although the latter topic has just started being discussed by the SEAMOBLY working group. They are not considered further here.

7.8. Context (or State) Transfer

This section starts with some examples of 'context' or 'state', before describing what the context transfer problem is and some possible solutions to it.

Examples of context (state) include:

- Header compression - The mobile and access router (AR) must have a synchronised, shared understanding of how the header is compressed over the air, so that it can be decompressed.
- Multicast group membership - The AR must know which multicast groups the mobile wants to receive.
- QoS policy - For example, the AR may have to police packets from the mobile, checking that they are within agreed limits and so protecting network resources.
- AAA profile - For example, security and the network's accounting policy for that mobile. IPsec state - The AR may act as an IPsec gateway, in which case, a security association between the mobile and AR enables packets to be encrypted and decrypted between the two.

Context is thus any information associated with some control protocol that affects how traffic is forwarded between the AR and a particular mobile. More generally, for some types of context, there may be information not just at the access router, but also at other routers on the data path (e.g. IntServ QoS). The wording below assumes the 'access router only' case; similar approaches should be possible for the general case.

Thus, the context transfer problem is to ensure that, when a MH moves to a new AR, the context is successfully regenerated at the new AR. Several ways have been suggested for doing this:

1. The mobile simply restarts the protocols after the handover. This may take some time, however.
2. The mobile updates the new AR with the state. However, at least for some types of state, this would require the AR to inform the mobile periodically about its state, so that the mobile has the correct information to upload when it moves.
3. The protocol responsible for the state is responsible for transferring its state. However, this requires modifications to these protocols, which the relevant IETF Working group may be reluctant to make, for example, if it is busy doing other things. Also, if a protocol is already widely installed, modifications will require substantial effort to roll out.

4. The state could be transferred by the handover protocol. This would require the handover protocol to be modified, and might be difficult if there is a lot of state to be transferred (larger than one packet, say).
5. Some central entity stores (a copy of) the state and downloads it to the new AR when required.
6. The old AR informs the new AR of its state when it knows that a handover is happening. A specialised protocol must be designed to do the job, one that can transfer whatever state is required.

It is not clear which is the best option, and at present, there is much active discussion at the IETF's SEAMOBY working group. Indeed, the answer probably depends on the particular protocol state and perhaps on the handover type (planned handovers appear to favour Option 6 more strongly than unplanned handovers). Options 4, 5, and 6 have the advantage that signaling is restricted to the wired network, so no extra traffic is sent over the air. A detailed analysis has just started. Some guesses:

- Option 1 is the simplest and may be best for most of the state where a seamless handover is not required.
- Option 2 does not seem viable in general, but may be feasible for some state (possibly multicast group membership?).
- Where a protocol modification appears simple, Option 3 may be best. For example, perhaps a handover could trigger an RSVP soft state refresh and so deliver the relevant state to the new AR.
- Since it is normally required that a handover is secured (e.g. to protect against a malicious user pretending it is an approved user handing over), it may be best for the mobility protocol to deal with (or at least assist) the transfer of essential security information. (Option 4 - for example in fast Mobile IPv6, it could be added on to the HI/HACK messages).
- Option 5 may be best when an obvious central entity exists, e.g. the MAP in hierarchical Mobile IPv6, perhaps to transfer AAA security credentials.
- Option 6 will be required if the other options are not possible.

Incidentally, the phrase 'context transfer' is often used to refer a specific new protocol for transferring control information from the old AR to the new AR (i.e. Option 6), as well as to the problem in general.

7.9 Paging and Dormant Mode Management

In a typical mobile system, at a particular moment, many mobiles are switched on but are not actively sending or receiving packets. Therefore, GSM and 3G networks (see [Chapter 2 & 3](#)), for example, invent a mode where the mobile is neither fully active nor off, but somewhere in between. The obvious idea is to do the same in a mobility-enabled IP network. Such a partially active mobile is said to be in dormant mode (alternative terms are idle and standby). There are benefits to both the mobile and the network:

- The mobile benefits because it can save its battery power. Typically, this is achieved by allowing the mobile to 'switch off' some of its power hungry

components during a sleep period. However, the network must be able to alert the mobile, for example, if a correspondent wants to communicate with it.

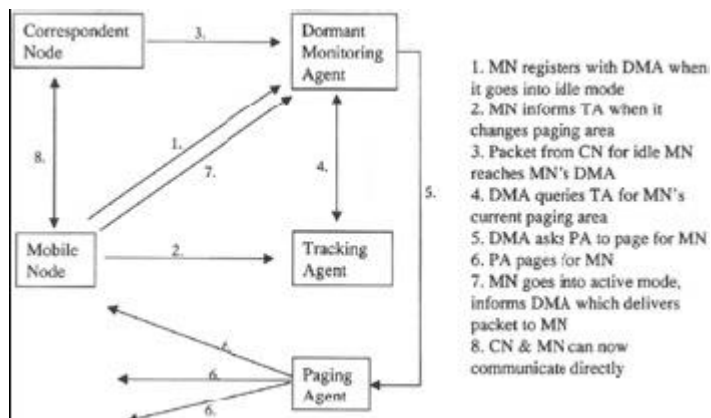
- The network benefits because the number of signaling messages over the air is reduced, and savings can be made on the routing state in the access network. The basic idea is to track a dormant mobile's location less accurately than an active mobile's location. This is done by defining paging areas (also called location areas), which consist of a number of access routers (ARs) corresponding to some geographical area and only tracking a mobile's location to the nearest paging area (rather than its nearest AR).

The process by which the network wakes up a dormant mobile is called 'paging'. Once a dormant mobile receives a page, it moves into active mode and informs the network of its exact location (i.e. its current AR).

A dormant mobile must thus be pageable. The method by which it achieves this, whilst still saving on battery power, is mainly a matter for the specific wireless link technology. A couple of options are that it periodically wakes up at well-defined times, or that there is a specific paging channel that is permanently monitored (whilst traffic channels are not monitored in dormant mode).

The final requirement is that a mobile must inform the network when it moves into a new paging area, and also occasionally remind the network where it is if it does not move. This process is called 'location updating', or 'paging area registration'.

The IETF's Seamoby working group has devised a functional architecture for dormant mode management, the key parts of which are shown in [Figure 7.1](#). In the actual physical implementation, the three types of agent could be merged.



7.9.1 Paging

The paging part of dormant mode management involves two main steps:

- To send the paging message to all the ARs in the paging area.
- To send the paging message over the air to the mobile.

7.9.2 Within the Access Network

Various paging schemes have been proposed. Most of these rely on multicast. Each paging area is identified by a multicast address, so a paging request sent to this address will propagate to all the ARs in the paging area (PA). For this job, multicast appears to be a good match to the requirements, since the paging message needs to be sent to several ARs and the multicast group membership will change only slowly (or not at all).

Some of the details differ between the various proposals, for example:

- The location information could be stored either centrally (say at the gateway) or distributed within the access network.
- The network could be told that the mobile is moving into idle mode either by an explicit message or implicitly; the latter could apply in a soft-state, per-host micromobility protocol, where the network can interpret the absence of routing refresh messages from a mobile as meaning that it has become dormant.
- The paging areas could be configured statically or dynamically, perhaps on a permobile basis according to its traffic and mobility pattern.

The differences between the various schemes are minor and reflect the origins of their particular underlying micromobility protocol.

7.9.3 Over the Air

The detailed issue here is the involvement of Layer 2 and 3 in sending a paging request to the mobile from the access router (AR). There are three options:

- Page at Layer 3 - This is the 'IP purist' approach. The procedure is (almost) independent of the radio technology, which makes it universal. However, it is difficult for the mobile ever to go into dormant mode, because it needs to be ready to receive a potential (Layer 3) page and to listen to (Layer 3) beacon messages from ARs (so it knows if it has moved into a new paging area).
- Page fully at Layer 2 - This is the 'traditional' cellular approach of GSM for instance. It is well optimised for power management purposes, because the mobile only needs to monitor the Layer 2 paging channel. But it may be less applicable if the AN includes several different link layer technologies.
- Layer 2 and 3 interact to achieve paging - One possibility is to insist that a Layer 2's broadcast channel can transmit IP packets (current channels cannot, though). Another possibility is to define a generic interface between the IP and a (generic) wireless link layer, with a primitive that causes a link layer page to be sent by the AR, and a corresponding one indicating its reception at the mobile.

7.10 A Brief Word on Security for Mobility Management

Security is a complex topic because it is about dealing with threats dreamt up by devious and clever people. This section looks very briefly at some of the threats introduced by wireless, mobile communications.

The first threat is that it is relatively easy for User A, the eavesdropper to listen in to User B's traffic over the wireless link - much easier than for a wired link, when User A must find the correct wire to tap, perhaps by breaking into User B's office. Many mobile systems use encryption over the wireless link to provide some degree of security - ranging from very good (e.g. GSM), to poor (e.g. the WEP algorithm for 802.11 wireless LANs is simple to crack).

The second threat is that User A can send a 'mobility update' message that pretends to come from User B - for example to tell the network that User B has moved to User A's IP address, so that User A steals User B's traffic. Such redirect attacks are probably the most important threat introduced by mobility. To overcome the threat, the techniques discussed (amongst others) are used. The basic method is:

- To establish a shared secret key between wherever the mobility updates will go. An obvious example is between a mobile host (MH) and its home agent (HA) in mobile IP. However, the key needs to be shared (distributed) in the first place. Although the MH can agree a key with its HA in advance, it cannot do so with its correspondents (for route optimisation). The general problem of key distribution is discussed briefly in [Chapter 1](#).
- To authenticate the 'mobility update' message. An algorithm calculates a digest that depends on the key and the fields of the message that need to be protected. When the network decrypts the digest, it checks that the key is correct. For example, in mobile IPv4, the default algorithm is HMAC-MD5.
- To protect against a replay attack, meaning that User A records a genuine (authenticated) 'mobility update' message and plays it back some time later. For example, in mobile IPv4, time stamps are used (so the home agent for instance checks that the time matches its own clock - within some error margin); it also allows nonces to be used (a 'nonce' is a random number that is used just once, which the mobile 'echoes' back to the home agent for instance).

Another threat is a denial of service (DoS) attack. For example, User A tries to block User B receiving any traffic by bombarding them with useless traffic. The general problem can be counteracted by firewalls for example. However, paging introduces a particular concern - can paging act as an 'amplifier' of User A's DoS packets? One idea would be for User A (A1) to send a packet to a fellow attacker, A2, who is in dormant mode. The page generates traffic in the network, but A2 maliciously does not reply with a location update. A1 tries repeatedly, but A2 never responds. Occasionally, A2 declares that it is still dormant, to stop the network assuming it is turned off (when it would not bother trying to page it). In fact it may be possible for User A to play the role of both A1 and A2.

A rather different threat associated with some mobility management schemes is that of location privacy. For example, route optimisation and per-host micromobility protocols might allow a correspondent to tell from a user's care-of address that they are away on holiday (and therefore easy to rob).

In general, it is very important that any mobility protocol is considered from a security point of view, to make sure that it does not have any obvious weaknesses that can be exploited and also that it does not open up a new security hole in another protocol.

7.11 Conclusions

This chapter has examined IP mobility. It has discussed personal and terminal mobility, and mobility solutions at the link, IP and 'application' layers. The focus has been on IP layer solutions to terminal mobility in an IP network that has a fixed infrastructure - that is, a fixed network with base stations that provide wireless connections to mobile terminals. Compared with link layer solutions, IP layer schemes should be easier to scale to large networks and also allow mobility management to be insulated from the details of the lower layers (future as well as current layers). This implies that the base stations should be IP routers (termed 'access routers'), with a plug-in card providing the particular air interface connectivity required to the mobiles.

IP terminal mobility has been broken down into two complementary parts - macromobility and micromobility - that is, mobility between Access Networks and mobility within an Access Network.

For macromobility, one option is the well-known Mobile IP protocol. Mobile IP (MIP) is the nearest thing to an agreed standard in IP-mobility, but its deployment has been very slow, mainly because of concerns about foreign agents (MIPv4) and continuing security issues (MIPv6). Another argument for MIP's slow progress (and that of IP mobility in general) is that its benefit is to allow a user to receive incoming calls - but this is not something that is currently wanted, since applications are either client-server based (e-mail, web browsing) or involve logging on to a server first (instant messaging). A critical test for MIP is its deployment in 3GPP2 (cdma2000).

Macromobility can also be dealt with by an application-layer solution such as SIP; although SIP has been designed for negotiating peer-to-peer sessions ([Chapter 7](#)), it also appears promising as an alternative to MIP.

Micromobility is now generally split into two problems - 'handover management' (i.e. making the actual handover as seamless as possible) and 'path updates' (i.e. ensuring that packets are delivered to the mobile terminal once it has finally moved on to the new access router (AR)). Both are under active discussion and research.

For 'handover management', particularly for planned handovers, the ideas of using a temporary tunnel between the old and new ARs, combined with handover smoothing (buffering and/or bicasting of packets), and the ability to transfer useful information (e.g. the mobile's prospective care-of address) look very promising.

For 'path updates', solutions fall into two classes. The first class, which is based on MIP, introduces local mobility agents that essentially act as a local proxy on behalf of the home agent. Such solutions should be relatively easy to roll-out, since only a few routers need to be upgraded with the special mobility agent functionality. The second class of solutions – perhost forwarding - involves a change to the routing protocol, i.e. to all the routers in the access network. This is (arguably) the 'IP purist' approach, since then the (single) IP routing layer is delivering packets to all terminals (fixed as well as mobile). The idea of having prefix-based routing on to which a host-specific routing 'tail' is added as the mobile node moves appears promising to improve the scalability of per-host forwarding schemes.

'Path update' protocols have been compared under several topics: operation, architecture, scalability, robustness, and philosophy. In general, protocols are much more similar to each other than is generally admitted. The solution favoured often depends as much on what someone believes an IP mobility protocol should do, as on technical details. This is partly because there is a lack of comparative modelling and implementation experience.

Currently, MIP-based solutions are being pursued in the mobileip working group at the IETF. Per-host forwarding schemes, however, have been moved from the IETF into the IRTF (which is the IETF's research arm), on the basis that further 'research' is needed before they are ready for 'engineering'.

Several other issues associated with IP mobility management have also been looked at briefly: context transfer and paging, which are being examined by the IETF's SEAMOBY working group, and security.

Finally, here are some predictions about the trends that will be important over the next few years in the field of IP mobility:

- IP mobility will become increasingly important with the rise of peer-to-peer communications on the IP networks or, to put it another way, as 'IP' plays an increasing role in '3G'.
- In the all-IP future, base stations will simply be IP routers, with a vast range of different link and physical layer technologies used to connect to mobile terminals. The access network will contain normal IP routers, which are enhanced with a little extra functionality to support mobility.
- However, new network architectures will become important, for example where users spontaneously form an extension to the network.
- The rise of SIP - particularly as it becomes ubiquitous as part of Release 5 of 3GPP and gets built into operating systems (e.g. Windows XP) - will lead to the widespread adoption of SIP- based macromobility management.
- Per-host forwarding solutions (for micromobility path updates) will make a comeback.
- IP mobility will be broken down into a number of independent protocols. This protocol separation is consistent with the IP design principles discussed in [Chapters 1](#) and [3](#), e.g. separate protocols for macromobility, handover management, path updates, paging, and context transfer.

- A standardised interface will be developed between the IP layer and a generic wireless link layer. For example, this will handle handover 'triggers' in a universal fashion (e.g. an indication that a mobile node is moving in range of a new access router).
- Increasing attention will be focused on aspects not traditionally considered part of IP mobility, because building an all-IP mobile network requires more than simply delivering packets to the correct access router.

Chapter 8

Quality of Service

8.1 Introduction

8.1.1 What is QoS?

The basic definition of QoS is given by the ITU-T in recommendation E.800 as "the collective effect of service" performance, which determines the degree of satisfaction of a user of a service.

There are a large number of issues, which affect user satisfaction with any network service. These include:

- How much does it cost?
- Can a user run the application they want?
- Can a user contact any other user they want?

None of these is a straightforward technical question. If a user wants to run a video-phone application, this requires that:

- The application is compatible with that used by the phoned party.
- The cost is not prohibitive.
- There is a network path available to the other party.
- The user does not have too many other applications running on their computer already, so that the computer has available resources.
- The network path can deliver all the required data packets in a timely fashion.
- The user knows the IP address of the terminal the other user is at.
- The end terminals can reassemble the data packets into a sensible order.
- The end terminals understand how to handle any errors in packets.

There are doubtless many other requirements. In identifying these requirements a few assumptions have already been made. In particular, the basic IP principles have been followed, as identified previously, and it has been assumed, for example, that much of QoS is a user/end-terminal responsibility.

Answering each of these questions leads to different fields of study within the general subject of QoS. These include:

- Traffic engineering - This includes how a network manager makes the most efficient use of their network, to reduce the cost.
- Policy management - Static network QoS provision, for example to give the boss the best network performance.
- QoS middleware - This is how a software writer creates generic components for managing both network and local resources so as to enable an application to be able to adapt to different situations.
- Control plane session management - As discussed in [Chapter 4](#), how users contact each other and arrange the most suitable session characteristics.

- Data plane session management - How end terminals make sense of the packets as they arrive.
- Network QoS mechanisms - How to build networks that can forward packets according to application requirements (e.g. fast).
- QoS signaling mechanisms - How networks and users communicate their QoS requirements.

Consideration of these last three issues, loosely described as 'User-driven Network QoS', is the focus of this chapter. The Internet today provides only one level of quality, best effort. It treats all users as equal. Introducing 'Quality of service' almost by definition means that some users, for example those not able to pay more, will see a lower QoS. Those who are prepared to pay more will be able to buy, for example, faster Web browsing. However, more importantly, introducing QoS also means that a wider range of applications will be supported. These include:

- Human - Human interactive applications like video-conferencing and voice.
- Business critical applications, such as Virtual Private Networks, where a public network is provisioned in such a way as to behave like a private network, whilst still gaining some cost advantages from being a shared network.

'User-driven Network QoS' is essentially about allowing users to request 'QoS' from the network. The type of QoS that may be requested could include:

- Guarantee that all packets for this session will be delivered within 200 ms, provided no more than 20 Mbit/s is sent.
- Guarantee that only 1% of packets will be errored, when measured over 1 month.

8.1.2 Why is QoS hard?

QoS, especially in the Internet, is proving hard to provide. QoS was actually included in the first versions of IP - the TOS bits in the IP packet header were designed to allow a user to indicate to the network the required QoS. Yet, to date, there is very little QoS support in the Internet. One of the problems appears to be in defining what is QoS and what the network should do questions that have been touched upon above. However, there are also a number of more pragmatic issues.

Cost/Complexity/Strength of QoS Compromise

Strong QoS can be obtained by giving each user much more capacity than they could ever use perhaps by giving each user a 100-Mbit switched Ethernet link. Clearly, this would be a very costly approach to QoS. Within the standard telephone network, high levels of QoS are achieved by placing restrictions on the type of applications that can be supported the telephone network only provides quality for a fixed data rate (typically 64 or 56 kbits). In a typical voice call, this resource is unused for more than half the time the caller is quiet while the other party speaks. A reasonable generalisation is that two out of the three are possible, e.g. low cost and low complexity lead to weak QoS.

QoS Co-operation

To achieve QoS requires co-operation between many different elements. One poorly performing network segment could destroy the QoS achieved. Similarly, an excellent network performance will not help the user perceived QoS if they are running so many applications simultaneously on their computer that they all run very slowly. An important question is what impact the nature of wireless and mobile networks has on QoS; to date, the development of IP QoS architectures and protocols has focused on the fixed Internet. This question is addressed extensively in this chapter.

8.2 Current IP QoS Mechanisms

Despite the fact that the need to provide QoS is a major issue in current Internet development, the Internet itself today does already provide some QoS support. The main elements in common use today are the Transmission Control Protocol (TCP), Explicit Congestion Notification (ECN) and the Real Time Protocol (RTP). This section reviews these mechanisms, with particular emphasis on their behaviour in a wireless network supporting mobile terminals.

8.2.1 TCP

The Transmission Control Protocol, TCP, is a well-known protocol that manages certain aspects of QoS, specifically loss and data corruption. It provides reliable data transport to the application layer. We will consider first how it provides this QoS service, and then consider the problems that wireless can present to the TCP service.

Basic TCP

TCP operates end to end at the transport layer. Data passed to the transport module are divided into segments, and each segment is given a TCP header and then passed to the IP module for onward transmission. The transport layer header is not then read until the packet reaches its destination. [Figure 8.1](#) shows the TCP header.

0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7
Source Port identifies the end point (application) of the connection								Destination Port identifies the end point of the connection								Sequence Number								Acknowledgment Number – next byte expected – previous bytes correctly received							
Header Length								Reserved								Flags								Receiver Window Size							
Checksum								Urgent pointer – if urgent flag set, this says where urgent data starts								Options (variable length)															

Figure 8.1: The TCP segment header.

The main elements of a TCP header for QoS control are the sequence number and checksum. When the TCP module receives a damaged segment, this can be identified through the checksum, and the damaged segments discarded. Data segments that are lost in the network are identified to the receiving module through the (missing) sequence numbers. In both cases, no acknowledgement of the data will be returned to the sender, so the data will be retransmitted after a timer expires at the sending node. The sequence numbers also enable the receiver to determine whether any segments have been duplicated, and they are used to order the incoming segments correctly. Receivers can provide flow control to the sender to prevent any receiver node buffer over-runs, by entering the 'window size', or maximum number of bytes, that the receiver can currently

handle. The sender must ensure that there is not so much data in transit at any one time that loss could occur through a receiver buffer overflow (Figure 8.2). To keep data flowing, receivers will send a minimum of TCP ACK messages to the sender, even if there is no data flow from receiver to sender.



Figure 8.2: Illustrating the sender's sliding window that limits congestion losses in both network and receiver.

TCP requires an initial start-up process that installs state in client and receiver about the transmission this state defines a virtual circuit. This state essentially identifies the two ends of the connection (IP address and TCP port identifier) and indicates the initial starting values for the sequence numbers. This is needed to ensure that repeat connections to the same destinations are correctly handled.

In addition to ensuring that its data rate does not cause a receiver buffer overflow, the sender is also responsible for preventing network router buffer overflow. TCP achieves this network congestion control by slowing down the data transmission rate when congestion is detected. This helps prevent data loss due to queue overflows in routers. To achieve this, the sender maintains a second window size that reflects the state of the network. The sender determines this window size by gradually increasing the number of segments that it sends (slow start sliding window protocol, Figure 8.3). Initially, the sender will send only one segment. If this is acknowledged before the timer expires, it will then send two segments. The congestion window grows exponentially until a timeout occurs, indicating congestion.




Figure 8.3: The size of the congestion window, which is the sender's understanding of the maximum data rate the network can support, grows. If the roundtrip time is large, or many timeouts occur due to loss, this growth is slow.

TCP requires neither network-based call admission control nor any support in routers, but it makes some assumptions about the nature of routers and transmission networks. In particular, this protocol assumes that transmission loss rates are small, so the overhead of end-to-end retransmission of corrupted packets is not an issue. It further assumes that loss

is primarily caused by network buffer overflows. It can be thought of as having an out-of-band, hard-state signaling protocol - the TCP handshake. The end terminals have traffic conditioning capabilities - they measure the traffic flow, and on identification of network problems, they can act on the traffic, essentially reducing the data transmission rate.

Wireless Implications for TCP QoS

Whilst the higher-level protocols should be independent of the link layer technology, TCP is typically highly optimised, based on assumptions about the nature of the link, which are not true in wireless networks.

The congestion control algorithm assumes specifically that losses and delays are caused by congestion. In a fixed network, if losses or delays are encountered, this implies a congested router. In this situation, the sender should reduce the level of congestion losses by slowing down its sending rate. This will reduce the required number of re-transmissions, thus giving more efficient use of the network whilst being fair to other users of the network. In a wireless network, losses occur all the time, independently from the data rate. Thus, slowing down does not alleviate the loss problem, and simply reduces the throughput. In a general network, there may be both wireless and fixed sections, and neither the sender nor receiver can know where losses have occurred and, therefore, what action should be taken on detection of losses. Since many wireless networks have a circuit-oriented link layer, any problems with TCP efficiency directly cause overall inefficient use of the link.

In the presence of frequent losses, the congestion avoidance algorithms have also been shown to produce throughputs inversely proportional to the round trip time. This is a problem as many wireless links have large latencies (as a result of the loss management required), and this problem would be compounded for mobile-to-mobile communications. Essentially, the reason for this is that the slow start algorithm and loss-recovery mechanisms both rely on the sender having data acknowledged - the time to obtain the acknowledgements depends upon the round-trip time. This result has been formally proven, but can be understood from [Figure 8.3](#), which illustrates the behaviour of the slow-start algorithm. This shows that, after a loss, the rate of growth of the congestion window (and hence the data throughput) is directly related to the round trip time. If losses are regular, this process will be repeated.

Whilst link layer error management, such as ARQ, can greatly improve the error rate of the wireless link, it achieves this with a cost of variable delay. TCP implementations use timers held by the sender to indicate when an ACK should have appeared in response to a transmission. If the timer expires, the sender knows to re-transmit the segment. Whilst the sender may attempt to set the timer based upon measurement of the round-trip time, it is difficult to do this accurately in a wireless network because of the random nature of the losses and ARQ-induced delays. It is possible, therefore, that the same segment is in transmission twice as the link layer attempts to send the segment across the link uncorrupted, whilst the sender has assumed that the packet is lost and has re-transmitted it. This is wasteful of bandwidth.

Another problem arises because wide-area modern wireless links typically have large latencies and large bandwidths. This means that at any particular time, a large amount of data could be in transit between the sender and receiver. If this value is larger than the receiver window, the sender will need to reduce its transmission rate, lowering the throughput, because the receiver buffer would otherwise run the risk of becoming overflowed. From the example given in RFC2757, for a UMTS network with a bandwidth of 384 kbit/s and a latency of 100 ms making the end-to-end latency 200 ms, the delay bandwidth product would be 76.8 kbits or 9.6 kbytes, compared with a typical receiver buffer or window of only 8 kbytes. Thus, unless TCP implementations are modified to have larger buffers, the data transmission will not fill the available capacity on the network - a terrible waste of expensive UMTS bandwidth.

Thus, to summarise the problems of TCP in wireless networks:

- Loss leads to a reduction of sending rate and so reduced throughput, but the loss remains as it was not caused by congestion.
- Loss leads to an initiation of the slow start mechanism. This is slowest to reach a steady state when round-trip times are large and will never reach a steady state if losses are frequent. This leads to reduced throughput.
- Variable delays lead to inaccurate time-outs, and so extra TCP re-transmissions will be generated, meaning that bandwidth is wasted on unnecessary re-transmissions.
- Large delays also mean that at any one time, a large amount of data will be in transit across the network. Therefore, the sender will have to suspend sending data until the data in transit have cleared the receiver's buffer.

Delay-related problems could be exacerbated on handover, which often increases the delay or even causes short breaks in communication. Ideally, TCP re-transmissions should be delayed during handover to allow the link layer to recover. However, techniques to provide this functionality and other solutions to TCP performance problems are still an area of active research.

A large number of solutions to these problems have been proposed. However, many produce knock-on effects. For example, if a TCP proxy is used as a proxy at the boundary between the fixed and wireless networks, the end-to-end semantics of TCP are broken. In addition to changing the semantics of TCP (the ACK does not mean the segment has been received), it then also breaks the IP level security model, and causes problems if the terminal moves away from that proxy. Other solutions may require changes to current TCP implementations, e.g. upgrades to every WWW server. Other ideas may require that a terminal uses different TCP implementations depending upon the physical network it is using for the connection a severe limitation for vertical handover. Finally, many solutions have been proposed that are not suitable because they may negatively affect the Internet stability, for example by leading to increased levels of bursts of traffic and congestion.

However, some modifications can be made. For example, the slow start process can be speeded up if the slow start initial window size is 2 or 3 segments rather than the traditional 1 segment. This has been accepted as a modification to TCP that does not affect the general stability of the Internet. Also, since slow start is particularly painful in wireless networks because of the long round-trip times, techniques should be used that

minimise its occurrence as a result of congestion losses. The use of SACK, Selective Acknowledgement, is also recommended. SACK is a technique that speeds up recovery where burst errors have damaged multiple segments - thus, its benefit depends upon the nature of the wireless network. It basically allows for multi- (rather than single) segment loss recovery in one round-trip time.

Whilst TCP proxies have many problems, application level proxies, however, may be of much greater benefit to the wireless environment especially as application layer protocols are often very inefficient. Even in this situation, however, the user must be in control of when and where proxies are used.

8.2.2 Random Early Detect and Explicit Congestion Notification

These are techniques that can be used by the network to reduce the amount of congestion losses, thus improving the quality of service. Random Early Detection (RED) has already been deployed within routers in some parts of the Internet. This technique deliberately discards packets as the queue builds up, providing a form of 'congestion ahead' notice to all users. Essentially, by dropping a few packets early on, it is possible to avoid congestion that would otherwise lead to larger numbers of packets being lost. Within the router, as the average queue length increases, the probability of a packet being dropped increases. Larger packet bursts will experience a larger packet-discard rate, and sustained loads further increase the packet-discard rates. Thus, TCP sessions with the largest open windows will have a higher probability of experiencing packet drop, causing them to start the congestion avoidance procedure.

Since the larger flows have a greater chance of experiencing packet drops, RED can avoid all the TCP flows becoming synchronised. This happens when the flows all experience congestion at the same time, all cut back, and all start to grow together. Explicit Congestion Notification is another mechanism to give advance warning of impending congestion. The router can mark, rather than just drop, packets with an explicit *Congestion Experienced* (CE) bit flag, on the assumption that the sender will see and react to this. In the case of TCP, the flag information must be echoed back by the receiver. Whilst ECN improves the performance of the network compared with packet drop RED, it requires changes to how TCP and IP operate and so, although it is now fully agreed within the IETF, it is unlikely to be introduced quickly.

8.2.3 RTP

The Real-time Transport Protocol, RTP, again provides end-to-end network transport functions. It provides ordering and timing information suitable for applications transmitting real-time data, such as audio, video, or data, over multicast or unicast network services. Again, we will first consider how it functions and then consider the impact that wireless networks could have on RTP.

Basic RTP

RTP requires no support in the network or routers. An initiation stage ensures that traffic descriptors are exchanged so that the end terminals can agree the most suitable traffic encodings. SIP is a suitable protocol for automating this stage. In addition to the RTP header information, RTP is usually run with RTCP, the Real Time Control Protocol. The

amount of control data is constrained to be at most 5% of the overall session traffic. RTP is a transport layer protocol that is typically run on top of UDP, extending the basic UDP multiplexing and checksum functionality.

RTP uses packet sequence numbers to ensure that packets are played out in the correct order. RTP headers carry timing information, which enables calculation of jitter. This helps receivers to obtain a smooth playback by suitable buffering strategies. Reception reports are used to manage excessive loss rates as, when high loss rates are detected, the encoding schemes for the data can be changed. For example, if loss is believed to be due to congestion, the bandwidth of transmission should be reduced. In other circumstances, redundant encoding schemes may provide increased tolerance to bit errors within a packet. This information can be delivered to the source through RTCP messages. The RTCP control messages provide information to enable streams from a single source, such as an audio and video stream, to be synchronised. Audio and video streams in a videoconference transmission are sent as separate RTP transmissions to allow low-bandwidth users to receive only part of the session. The streams are synchronised at the receiver through use of the timing information carried in the RTCP messages and the time stamps in the actual RTP headers. Full stream synchronisation between multiple sources and destinations requires that sources and receivers have timestamps that are synchronised, for example through the use of the network time protocol (NTP).

To prevent interaction between RTP and the lower layers, application frames are typically fragmented at the application level - thus, one RTP packet should map directly into one IP packet. RTP provides a means to manage packet re-ordering, jitter, and stream synchronisation, and can adapt to different levels of loss. However, it cannot in itself ensure timely delivery of packets to the terminal. This is because it has no control over how long the network takes to process each packet. If real-time delivery or correct data delivery is required, other mechanisms must be used.

Mobility Issues for RTP QoS

While RTP is largely independent of mobility, the overall RTP architecture includes elements such as mixer and translator nodes for service scalability and flexibility. If the core network includes several of these components, the mobility of the terminal may lead to situations where the mixer and the translator may change. These nodes have been pragmatically introduced as a means to handle multicast sessions. In large sessions, it may not be possible for all users to agree on a common data format - for example, if one user has a very low bandwidth link and all other users want high-quality audio. Mixers, placed just before the start of a low-bandwidth network can be used to overcome some of these limitations by re-coding speech and multiplexing all the different audio streams into one single stream, for example. This is done in such a way that the receiver can still identify the source of each element of speech. Translators are used in RTP to overcome some problems caused by firewalls.

Wireless Issues for RTP QoS

Low Battery Power

RTP makes large use of timing information to achieve its full functionality. The clocks used for this need to be synchronised across the network. The Network Time Protocol,

NTP is typically used for this. However, for NTP to provide the required high levels of accuracy (approximately in the microsecond range) it could require that the mobile terminal has IP connectivity for significant time periods (hours or days). This is somewhat unrealistic given current battery lifetimes. Therefore, some alternative mechanism to allow quicker convergence to NTP may be useful for mobile nodes. If the base stations were high-level NTP servers, it is possible that good synchronisation could be maintained here, which would enable much quicker convergence for the mobile terminals - however, this is a requirement (albeit simple) on mobile networks to provide this additional service to their users.

Compressible Flows

For low-bandwidth links, the header overhead of an RTP packet (40 bytes) is often large compared with the data - this is particularly important for Voice over IP traffic (20 bytes of data per packet for a voice packet encoded at 8 kbit/s, packets every 20 ms). In these circumstances, RTP header compression is often used. This is operated on a link-by-link basis. It enables the combined RTP/UDP/IP header to be reduced from 40 bytes to 2 bytes. No information needs to be transmitted to the link layer to achieve this compression. Because the RTP compression is lossless, it may be applied to every UDP packet, without any concern for data corruption. To save processing, as it is likely that the only traffic that will benefit is RTP, heuristics could be used to determine whether or not the packet is an RTP packet - no harm is done if the heuristic gives the wrong answer. For example, only UDP packets with even port numbers should be processed (RTP always uses an even port number, and the associated RTCP uses the next, odd, port number), and records should be kept of the identity of packets that have failed to compress.

However, this process only works once the data are being transmitted. If the application wants to improve QoS by reserving resources within the network, the application does not know if link-layer compression will be used, and the network layer does not know that compressible data will be transmitted. Thus, an application will request a reservation for the full data bandwidth. This reservation may be refused over the wireless link because of insufficient bandwidth, yet the compressed flow could be easily served. Without passing application layer information to the link layer, the link layer will need to manage this possibility intelligently.

There are two options:

- Allocate for the full bandwidth request initially, but reduce the local link-layer reservation on detection of compressible (usually RTP) traffic. Although this may lead to reservations being refused unnecessarily, it would allow the unused portion of a reservation to be recovered.
- Assume that RTP is used for all delay-sensitive reservation classes, and under-allocate bandwidth accordingly. Since the vast majority of real-time traffic will use RTP, this may be a suitable solution - although the traffic will need to be monitored to detect and correct when this assumption fails.

For all transmissions, not just RTP transmissions, the overhead of the IP packet header can be reduced. A number of header compression schemes do exist, particularly if the underlying link appears as a PPP, point-to-point protocol, links to the IP layer above.

However, TCP or RTP header compression is incompatible with network layer encryption techniques. Another possible problem with compression is that even a single bit error in a compressed header could lead to the loss of the entire segment for TCP, this would lead to the slow start process being triggered. It is assumed that payload compression will take place at the sending nodes, in an effort to reduce the cost to the user of the link (assuming that cost to the user is directly related to bandwidth used).

8.2.4 Conclusions

A limited set of basic QoS functions is already available within the Internet. However, none of these mechanisms can support real-time services, as they cannot ensure timely packet delivery. To do this would require some support by the network itself - the network will need to be responsible for more than just attempted packet delivery. This has been an active research area within the IETF over the last few years, and indeed, some progress has been made over the last year towards introducing QoS into IP networks. This problem is examined in the next sections of this chapter.

Further, to date, much Internet development has ignored the problems that mobility and wireless could cause. This is also true of many of the newer IETF standards. Although this situation is rapidly changing, some of the problems are fundamental, as to overcome they would require changes to the huge installed base of TCP/IP equipment, so many of the issues are likely to remain for many years. To some extent, it means that innovative solutions to minimise the impact of these problems need to be provided by the wireless link layers. This may be one area in which wireless network solutions may differentiate themselves.

8.3 Key Elements of a QoS Mechanism

QoS is a large topic and, as previously indicated, has implications in every part of the system. The first stage in understanding the problem is therefore to attempt to structure the QoS problem into smaller units. This section identifies what the basic elements are, and looks at some of the different design choices that exist for each of the elements.

As part of this, the problem that needs to be considered is: What is the required functionality within the network to provide QoS? The mechanisms that can exist within the routers, to enable the network to provide QoS, are examined later in this chapter. Since network QoS is essentially about giving preferential treatment to some traffic, there need to be mechanisms to negotiate this treatment with the network. This process is covered under a discussion of signaling. Finally, mechanisms are needed that allow the network to ensure that admission to the service is controlled - after all, not every user can have preferential treatment at the same time. Throughout this section, special attention is paid to issues caused by wireless and mobile networks.

8.3.1 Functionality Required of the Network to Support QoS

Quality of service may require control of a range of features including packet delay, packet loss and packet errors, and jitter. Beyond a basic minimum, QoS is meaningful to a user only if it exists on an end-to-end basis. As an example, the error rate of the data, as finally delivered to the application, is more significant than the error rate at any particular point within the network. As previously discussed, many aspects of QoS, including packet loss, stream synchronisation, and jitter, can be controlled at the terminal through

the use of suitable end-to-end layer protocols. As an example, the transmission layer protocol TCP is currently used to control error rates, whilst RTP is used to manage jitter and stream synchronisation. The only parameter that cannot be controlled in such a fashion is the (maximum) delay experienced by a packet across the network [1]. Providing delay-sensitive packet delivery requires co-operation from each element within the network. This leads to a division of responsibility for QoS according to Figure 8.4.

Layer	Example Protocols	Functionality Required
Application		
Transport	TCP, UDP, RTP	Error control, Stream Synchronisation, Jitter Control
Network	DiffServ	Timely Packet delivery
Link	FEC	Timely frame delivery, some error management, orderly frame delivery

Figure 8.4: Internet Layer Model with QoS protocols and functionality.

Whatever functionality is placed within the network to support QoS, this functionality, or its effects, needs to be clearly described to users. In general terms, users can be easily bewildered by a totally flexible system. It may be possible to offer a huge range of services, each with different probabilities of being maintained. However, as described in Chapter 2, UMTS networks define only four classes:

- Conversational - For applications such as video telephony.
- Streaming - For applications such as concert broadcast.
- Interactive - For applications such as interactive chat, WWW browsing.
- Background - For applications such as FTP downloads.

However, it has been proposed that even these classes could be collapsed into only two delay sensitive and delay insensitive as evidence exists, which suggests that only two classes can be truly distinguished within the Internet.

Finally, it is worth stating that just because it is implied here that only delay is important, this does not necessarily mean that only delay will be controlled by the delay-sensitive class. Jitter may be controlled as part of this, either explicitly, or by controlling the maximum delay that traffic experiences. Some effort may also take place to prevent congestion losses in such a class.

8.3.2 Interaction with the Wireless Link Layer

Although, above, a picture has been drawn with clear separation between layers and functions, life is never so clean-cut, and interactions will exist between different elements. These interactions are most obvious - and most problematic - between the network and link layer. Network layer quality typically manages the router behaviour in order to achieve a certain quality of service. This works well if the link is well behaved if it has no significant impact on the delay, jitter, or loss that a packet might experience. However, this is not true with a wireless link layer. Furthermore, the simplest method to overcome these problems bandwidth over-provision is not practical in general in a wireless environment, as bandwidth is expensive. For example, in the recent UK UMTS spectrum auction, 20 MHz of bandwidth went for 4 billion UK pounds. Therefore, link-

layer mechanisms are needed to manage the quality of data transmission across the wireless network. It is important that any quality provided at the network layer is not compromised by the link-layer behaviour. This could occur, for example if the link layer provides some form of re-transmission-based loss management, (such as ARQ) without the network layer being able to control the delay, or if the link layer reorders packets that have been sent for transmission. The [next section](#) expands upon these issues that have a significant impact on QoS.

Loss Management

There are a number of problems that wireless networks have that lead to data loss.

Low signal-to-noise ratio

Because base stations are obtrusive and cost money, they are used as sparingly as possible, and that means that at least some links in every cell have very low signal-to-noise ratios, and thus very high intrinsic error rates. Typically, a radio link would have a bit error rate (BER) of 10^{-3} compared with a fiber link with a BER of 10^{-9} .

Radio Errors Come in Bursts

In many other networks, errors are not correlated in any way.

Radio Links Suffer both Fast and Slow Fading

Fast fading causes the received power, and hence the error rate, to fluctuate as the receiver is moved on a scale comparable with the wavelength of the signal. It is caused by different rays traveling to the receiver via a number of reflections that alter their relative phase (a GSM signal would have a wavelength of 10–20 cm or so). Slow fading - also called shadowing is caused by buildings and typically extends much further than a wavelength. There are solutions to these problems. To overcome the high error rates, radio links employ both forward and backward error correction. For example carrying redundant bits enables the receiver to reconstruct the original data packet (Forward Error Correction), whereas Automatic Repeat Request (ARQ) is used to re-transmit lost packets. Where ARQ can be used, the error rates can be reduced sufficiently such that any losses TCP sees are only the expected congestion losses. However, this scheme relies on link-layer re-transmissions and so significantly increases the latency, which can be a problem for real-time traffic. To counter the problem of burst errors and fast fading, radio systems can mix up the bits and fragment IP packets into smaller frames. Again, this could cause latency problems for real time traffic.

All these techniques, however, still do not take into account the unpredictable and uncontrollable errors that might occur. An example of such a problem could be when a user of a wireless LAN moves a filing cabinet, or a user of a GSM system is on a train that enters a tunnel. In such situations, the signal level might fall so far into the noise that the session is effectively lost. Mechanisms also exist so that the wireless transmitters can control to some extent how the errors appear to the terminal. For example, some traffic (such as voice) prefers an even selection of bit errors to whole packet losses. Certain encoding of video traffic, however, leads to a preference that certain whole video packets are dropped, ensuring that the top priority packets are transmitted uncorrupted. If the link layer knows the type of traffic carried, it can control the loss environment, by using different error correction schemes. However, this requires significant communications

between the application and the wireless layer. Exchange of wireless specific information from the application is generally considered a bad thing. It breaks the design principles described in [Chapters 1 and 3](#), which state that the interaction between the link layer and the higher layers should be minimised. Higher layers should not communicate with the link layers, and protocols should not be designed to the requirements or capabilities of the link layer. Applications and transport layer protocols should not have 'wireless aware' releases.

So, how can these issues to be handled? The error rate on wireless links is so bad that it is a fair assumption that error correction techniques should be used wherever possible. It is assumed that forward error correction is always used on the wireless links to improve the bit error rates. Ideally, for non-real time services, the errors on a link should be controlled to produce an overall error of no more than 1 in 10⁶. For real-time service, the errors should be corrected as much as possible within the delay budget. This implies some mechanism for communicating QoS requirements down to the link layers, perhaps using the IP2W interface. Furthermore, to enable wireless loss management techniques to be used, network providers should assume that a significant proportion of any delay budget should be reserved for use within a wireless link.

Scheduler Interactions

Once QoS exists at both the link and network layers, there is a possibility for interactions between the two QoS mechanisms. There is unlikely to be a one-to-one mapping between network and link-layer flows. So, in the general case, thousands of network layer flows may co-exist, whereas there is usually a limit on the number of physical flows that may exist. In the general case, there cannot be a one-to-one mapping between these flows and the queues that are used to put these flows on to the network. With multiple queues at both layers, there will also be scheduling to determine how to serve these queues at both layers. This can cause problems. Consider a simple case where the network has a 'fast' and 'slow' queue for IP packets. The network layer bounds the size of the fast queue (to say 50% of the available bandwidth) to ensure that the slow queue is not starved. If there is a packet in both queues, the fast packet will always be delivered to the link layer before the slow packet. The link layer also has two queues: 'first transmission attempt' and 're-transmission'. These queue link-layer frames (parts of IP packets) and are served in such a way that frames for re-transmission are given a slightly lower priority than 'first attempt' frames. Now, suppose the IP layer sends first a 'fast' packet, which is divided into two frames at the link layer, and then a large TCP packet.

The second half of the fast packet fails to be correctly delivered, is queued for retransmission, and is then blocked for a long time as the large TCP packet is served from the 'first attempt' queue. Although this is a simplistic scenario, it illustrates the points that:

- The network layer needs a clear understanding of the behaviour of link-layer classes and link-layer scheduling to prevent interactions between the behaviours of the two schedulers.
- The link layer needs QoS classes that support the network requirements.

Again, this implies some mechanism for communicating QoS requirements down to the link layers.

8.3.3 Mechanisms to Provide Network QoS

When traffic enters a router, the router first determines the relevant output for that traffic and then puts the packet into a queue ready for transmission on to that output link. Traditionally, in routers, traffic is taken (scheduled) from these output queues in a first come, first served basis. Thus, as illustrated in [Figure 8.5](#), packets can be delayed if there is a large queue. Packets can be lost if the queue is filled to overflowing.




Figure 8.5: Normal router behaviour leads to uncontrollable delays and packet losses.

QoS implies some kind of preferential treatment of traffic in routers. This preferential treatment may be allocated on a per-flow or aggregate basis. A flow is an associated sequence of packets flowing between the same source/destination pair such as a sequence of packets that make up a voice transmission. Individual flows can be aggregated together into a shared class. Per-flow scheduling gives better control of the QoS, enabling firm guarantees to be made about the treatment of the traffic. However, this also requires that per-flow state be maintained in every router, and this per-flow state is used to determine the scheduling of packets through the router. This causes scalability problems within the core network, where large numbers of individual flows may co-exist. In the alternative aggregate treatment, traffic on entry to the network is placed into one of a few traffic classes. All traffic in any class is given the same scheduling treatment. This solution gives better scalability and can also reduce the complexity of scheduling algorithms. In the general case, however, less firm guarantees can be given to a user about the nature of QoS that they can expect to receive as illustrated in [Figure 8.6](#).




Figure 8.6: Aggregate scheduling gives less predictable behaviour than per-flow scheduling.

In certain cases, it is possible to use traffic aggregates for scheduling whilst still achieving hard QoS guarantees on a per-flow basis, and one such example is discussed later. In general, however, such techniques can only provide hard guarantees for a small number of QoS parameters.

Thus, we can see that the type of QoS functionality that we wish to provide has a direct impact upon how easily it can be supported by routers. Broadly speaking, simple QoS services can be supported by simpler scheduler implementations. More complex QoS services with many parameters to be controlled may require very complex techniques for managing the queues within the routers.

When QoS is used at the network layer, once the traffic reaches the first router, it is scheduled in order to achieve the required service. However, in the wireless world, huge problems could occur in sending the data to the first router. Thus, there needs to be a link-layer mechanism that ensures that QoS is controlled across the first link into the Internet. This QoS protocol is link-layer-specific. It is assumed that the IP module understands the mapping between the QoS protocols that it supports and the link layer protocols and QoS classes. For example, the IP module may actually use some form of link-layer reservations for the top-priority prioritisation traffic.

8.3.4 Signalling Techniques

Prioritization and Reservation

There are two main types of QoS solutions reservation-based solutions and prioritization based solutions. They essentially differ in how the user communicates to the network about their requirements. In reservation-based services, a node will signal its request for a particular QoS class prior to data transmission. By providing a description of the data traffic, it is possible for the network to use this information to allocate resources, on either a per-flow or aggregate basis. This enables a large degree of control over the use of the network, and hence can provide a good degree of confidence that the required quality will be achievable.

In contrast, no advance network negotiation takes place with prioritisation services. Traffic is simply marked to indicate the required quality class and then sent into the network. This type of system can only ensure that higher-priority traffic receives a better quality of service than lower-priority traffic. In most practical implementations, this will be augmented by 'service level agreements'. These contracts may be thought of as a static signaling mechanism. They may be used to restrict the amount of top-priority traffic transmitted from any particular source, enabling the network provider to make stronger probabilistic assurances of the nature of service that top priority traffic will receive.

Characteristics of Signalling Systems

To enable efficient use of scarce resources whilst also maintaining strong probabilistic service guarantees, it is assumed that, especially in the 3G environment, some reservation signaling will be required for real-time services. The signaling may be carried with the data, which is known as in-band signaling, or it may be out-of-band and thus separate from the data. Inband signaling ensures that the information is always carried to each router that the data visit, which is useful when routes change frequently, as in mobile networks. Out-of-band signaling, as used in telephone networks, is more easily

transported to devices not directly involved in data transmission - for example admission control servers. Most importantly, however, is the fact that in-band signaling requires an overhead to be carried in every data packet. A simple in-band signaling system, requesting only real-time service for a specified bandwidth of traffic, could add an approximate 10% overhead to a voice packet. The signaling may be soft state, in which case, the reservation needs to be explicitly refreshed. This makes it resilient to node failures. Conversely, a hard-state protocol can minimise the amount of signaling. The telephone network uses hard-state signaling the caller dials only once. With a hard-state signaling protocol, care needs to be taken to correctly remove reservations that are no longer required.

Different models exist in terms of the responsibility for generation of the signaling messages. These models are often coupled with responsibility for payment for use of the network. In a UMTS style of network, the mobile node is responsible for establishing (and paying for) the required Quality of Service through the mobile network domain for both outbound and inbound traffic. This model does not require that both ends of the communication share the same understanding of QoS signaling. It is a useful solution to providing QoS in a bottleneck wireless network region. The mobile user essentially pays for the privilege of using scarce mobile network resources. However, it is less easy to provide true end-to-end QoS in this situation. Inter-working units need to exist at each domain boundary that map between different QoS signaling and provisioning systems, and this inter-working may break IP security models. This solution typically assumes that the inbound and outbound data paths are symmetric true in the circuit-switched networks in which this model was developed, but not necessarily true in the IP packet network. Other solutions have one party responsible for establishing the QoS over the entire end-to-end path. The standard Internet models assume that the receiver is usually responsible for QoS establishment, as they receive value from receiving the data. However, these solutions usually require that the data sender also participate in any signaling and they retain ultimate responsibility for any payment this is seen as a possible mechanism for limiting 'junk mail'.

Wireless Efficiency

The limited, expensive wireless bandwidths mean that great efforts are required to minimize any signaling overhead carried on the link. This implies the need for hard-state signaling over the wire. This is easily achieved using RSVP (discussed later), which allows the softstate period to be set on a hop-by-hop basis, although additional functionality is then required to protect the network against hanging reservations - reservations left as a result of incorrect application termination. Further optimisation of signaling attempts to use one signaling message for several purposes. As an example, the link-layer QoS request could be designed to contain enough information (including destination IP address) to enable the wireless access router to establish the network QoS without the need for the mobile to transmit a network layer message. Avoiding this type of protocol coupling/layer merging helps to protect the future flexibility of the network. Similarly, improved efficiency implies the need for wireless application specific information to be passed to the wireless access router. However, unless wireless specific applications are to be developed, for example using RSVP with wireless extensions, such information would need to be configured into a link-layer driver.

8.3.5 Admission Control

Call Admission Control Architectures

QoS treats some traffic preferentially to others, and this implies the ability to reject traffic. The call admission functionality may exist at various places within the network. These alternatives are illustrated in Figure 8.7.




Figure 8.7: Different locations for the call admission functionality in different subnets.

In some solutions, each router is responsible for their own call admission decisions. Coupled with the Internet resilient routing, this enables a system that has no single point of failure. Each node makes a decision based on its complete knowledge of its current (local) state. However, such a solution does not scale well. The processing overhead of call admission would be significant for core routers, which route many thousands of flows.

Therefore, an alternative solution is that only edge routers process-call admission requests. These nodes use their local knowledge of their current state and make some assumptions about the rest of the network that enables them to make a decision on behalf of the core of the network. In particular, they can assume that no traffic enters the domain without being subjected to the same call admission scheme. The statistical effects of large numbers of flows facilitate the decision making process.

A fully centralised system enables a decision to be based on global as well as local knowledge. An edge router directs the request to the centralised admission control unit. There are a number of benefits of this approach. In addition to avoiding the scalability problems of the hop-by-hop approach, this scheme may enable QoS where routers have no QoS support. It allows the call admission mechanism to be upgraded and replaced easily. Centralised admission is not well suited for schemes with per-flow routing, as then, large amounts of communication would need to exist between every router and the call admission server. In addition, certain types of call admission criteria, in particular delay-based admission, are less well suited to centralised admission schemes. This is because the centralised unit can never know the actual full state of the network. Similarly, centralised admission schemes are less suited to the mobile environment where the state of the network is likely to change rapidly.

Admission Control Descriptions

Call admission may be based on a number of parameters that describe the traffic. Increasing the number of parameters enables more accurate admission decisions, leading to more efficient network usage. However, such decisions tend to be more complex and require a full analysis of the traffic characteristics of each existing flow. At the other extreme, the information offered could be simply the maximum bandwidth required. Such a minimal approach reduces the amount of information that needs to be signalled across a network. This approach is also more suitable when centralised call admission needs to be supported. This is because a centralised unit can only ever have an approximate understanding of the state of the network. Therefore, to make admission control choices, it needs to use approximations that have been found to be most possible if bandwidth-based admission is used. From the point of view of a network operator, the use of peak bandwidth as the main traffic descriptor also simplifies billing - as the operator can make a bill based on that one parameter that reflects simply how much actual network resources are used. A user can minimise their bill by doing traffic shaping to keep the required peak bandwidth as low as possible.

Traffic Classification and Conditioning

Once the network has accepted that the data can be transmitted, and the data are actually being transmitted, there are a number of functions that need to be provided to ensure that the network is protected against malicious use. As in call admission, these functions may be provided on a hop-by-hop basis, or solely on entry and exit to a network. By using these functions on exit from a network (and terminal), steps can be taken to ensure that transmitted data are within the contract, so that the behaviour through the network is understood. Throughout this section, the term 'QoS contract' is used to refer to either the dynamically signalled QoS contract or the static contract described through the service level agreement. The first stage, classification, is to identify the flow to which traffic belongs, through analysis of the packet header. The packet can then be associated with a particular QoS contract.

As an example, packets between a particular source-destination pair may be entitled to real-time forwarding provided that a strict bandwidth limit is not exceeded. Once the stream has been identified, traffic meters measure its temporal properties against the QoS contract. This meter may pass state information to other conditioning functions to trigger specific actions for each packet that is either in- or out-of-profile. One action that may be triggered by the measurement is traffic shaping. This is the process of delaying packets within a traffic stream to cause it to conform to the traffic profile. A shaper may be used, for example, to smooth out the burstiness of traffic, as traffic that is near constant bit rate can be managed more easily. This process in particular might take place in a terminal. A packet marker might be used to label traffic to ensure that it receives the required QoS treatment through that network domain. Additionally, packet markers might change the marking on a packet if the traffic has broken the agreed contract. This re-marking ensures that the traffic does not damage the QoS of in-contract traffic by ensuring that out-of-contract traffic does not receive the requested QoS. This re-marking also acts as a signal to the receiver that the QoS contract was violated - enabling action to be taken by the

end-to-end application. Packet droppers, which simply drop packets, provide another means to handle traffic that has broken the agreed contract.

Mobility Issues

The call admission architecture used has a significant impact on the problems that might be experienced during handover. Therefore some of these issues are examined. Solutions to overcome these problems are in turn affected by the nature of the router and signaling mechanisms used to provide QoS.

QoS Management during Handover

A handover occurs when a mobile node changes its point of attachment to the network. This implies that the route taken by data will change. Any QoS that has been established for that data, and particularly any reservation, will therefore be disrupted. To ensure minimal disruption during handover, a number of alternative mechanisms could be used. These are discussed in order of increasing complexity. For prioritisation QoS, little needs to be done to manage QoS during or after handover, as all class descriptions are relative and assurances are statistical.

The problem is more complex for reservation-based QoS (Figure 8.8), where some service guarantees have been made. A reservation-based handover is described as seamless if the application or end user cannot identify that a mobility event has taken place. To some extent, this could be managed through careful descriptions of the service classes - for example, by stating that traffic will be delivered within a certain time bound only 90% of the time. Thus, no attempt is made to provide QoS during handover, simply to re-establish the QoS reservation after handover has taken place.




Figure 8.8: Illustrating how handover can affect reservation based QoS.

One improvement is that each node simply reserves a portion of its available bandwidth to be used solely for traffic that enters the node as a result of handover. This is known as a 'static guard band'. There needs to be a mechanism to enable nodes to identify the handover traffic, and also the requirements of that traffic. This mechanism needs to be quick, as packets affected will already be in flight. One way to manage this exists if an aggregate, class-based service is provided, and packets carry an explicit QoS class marking in the IP packet header (as is provided, for example, in the aggregate DiffServ approach to QoS). Handover traffic can then be re-marked to the (network internal) handover class associated with the reservation-based service identified by the original class marking. For each reservation class, there exists a guard band for use by this handover marked traffic. Traffic should be remarked into this class by a node that recognises that a route change has occurred; this node will assume that the route change has been caused by handover. This assumes that state information about this reservation is held at any node at which the data path might diverge.

Otherwise, there is no way of identifying that a route change has occurred for that particular traffic flow. Thus, if the route were to diverge at any point within the domain (as might occur if a per-host routing mobility management scheme is used), every node within that domain must have per-flow state. This is illustrated in Figure 8.9. If a tunnel-based mobility management scheme is used, the tunnel anchor nodes will need per-flow state.

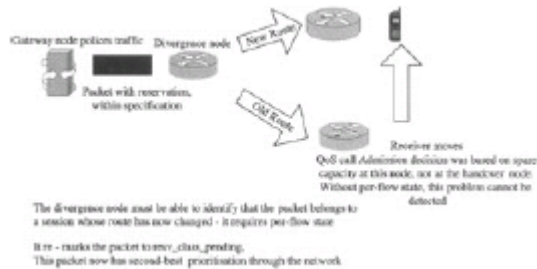


Figure 8.9: Per-flow state is required at routers if reservation based traffic is to be easily identified during the handover process.

Thus, we can see that reservation-based QoS cannot easily be used with the simplest edge admission control schemes. To do so would severely weaken the strength of QoS guarantees that could be made, as the amount of traffic in any one class cannot be limited to any particular node, because paths through the network could change rapidly as a result of mobility after the reservation process has been completed.

The use of static guard bands means that bandwidth may not be used efficiently, or more complex router schedulers will be needed to enable best-effort traffic to use the guard band when not required by handover traffic. This problem would become worse if a large number of different reservation classes were supported. This system can be improved upon when global information can be used to make the admission control decision - for example, using a centralised network management system - as then the size of this guard band can be adjusted dynamically according to the traffic distribution around the network. Such a node may instruct routers to reserve a larger guard band if the neighbouring cells are all busy, as handover is more likely to occur in this situation. Where nearby cells are empty, very little capacities needs to be allocated for handover traffic. The nature of these policies, their complexity, and the assumptions they make about user mobility are all areas under current research.

Using such procedures, it is possible to design a network such that there is a high probability that, if the new route can support the handover, the handover itself will be seamless. However, it is also possible that, for example, six sessions simultaneously handover, of which five can be supported through reservation once handover is completed, but during handover, all six suffer degradation the guard band, becomes too full. To avoid this situation, either the nodes involved in handover must communicate QoS state, which couples the mobility and QoS protocols, or the admission control must not be based on purely local decisions. An alternative approach is required for traffic that does not carry a QoS class marking, as is typical for traffic that is processed on a per-flow basis as in Integrated Services. In this situation, each session could make reservations for itself in nearby cells in preparation for likely handover. Thus, the session inserts per-flow

state at a number of nodes within the network. (These reservations could be considered a form of dynamic guard band.) Since many reservations may be needed, although only one will be used, this again could waste bandwidth. Thus, these reservations are 'passive' the space is used by best-effort traffic until the reservation is made 'active'. Since the mobile node does not, and should not, know the network topology, pre-reservation is improved by making the base stations responsible for the passive reservations rather than the mobile node. This also reduces the signaling load on the mobile node. Such pre-reservation schemes are difficult, as the route through the network is not usually identified until handover has taken place. Therefore, such schemes couple the mobility management and QoS reservation process. For example, some propose that this passive reservation be made between every probable handover cell and a mobile IP home agent. All traffic in both directions is forced to flow through the home agent this removes the possibility of any route optimisation. Other approaches propose that the nearby base stations do passive reservations on their wirelink, and that the current base station performs a reservation to each of the nearby base stations. Then, all traffic is forced to go through to the first base station, which then forwards the traffic.

Context Transfer Protocol

When handover occurs, network layer QoS typically needs to be re-established. In a wireless network, the weakest link is often the wireless section. It is most important that the reservation is established here extremely quickly; even if no action is taken to manage the network QoS until after handover is completed. It is assumed that the Layer 2 reservations will be established as part of the handover procedure. However, Layer 3 information is required if the wireless access router is to be able to associate incoming IP packets with the relevant link-layer reservations. We assume here that there is a context transfer protocol that transfers such information between the wireless access routers, as illustrated in [Figure 8.10](#).




Figure 8.10: Use of context transfer protocol.

The context transfer protocol does not yet exist. It is being developed by the SEAMOBY working group within the IETF in order to speed up the handover process. The context transfer protocol might include security information or QoS state information.

QoS Management after Handover

Where passive reservations have been used to manage the handover process, there is no need for any further QoS restoration after handover has occurred. Essentially, the call admission process was carried out in advance of the handover. The other approaches to handover management, however, require that the QoS reservation is formally re-established after handover. Ideally, this process should be confined to the region

impacted by mobility and should not involve the mobile node in order to conserve battery power and wireless bandwidth. Later, a discussion of RSVP, a specific QoS signaling mechanism, will show how this can be achieved.

When the mobility management mechanism relies on establishing tunnels, it is the responsibility of the tunnel end points (the mobility agent and the wireless access router) to re-establish the QoS. For efficiency and ease of processing, tunnels typically support a large number of flows within the one tunnel, but this would require that the tunnel be established with QoS suitable to support the highest QoS flow, which would usually be wasteful of resources. Otherwise, multiple tunnels need to be established, for example, one for each QoS class. Tunnels also need to be able to correctly route all control messages so reverse tunnels will be needed in the case of RSVP. Whilst a number of Internet Drafts and RFCs exist addressing the problems of tunnels and QoS, it is clear that this issue presents a large number of practical difficulties, with additional processing requirements and restrictions on network topology.

In turn, this actually also constrains the maximum jitter that a packet may experience
Maximum jitter = maximum delay-fixed transmission delay.

Easily justifiable if one thinks of a congestion loss as a packet with infinite delay. Other than the transmission delay, the time it takes to transmit bits from one end of a cable to another is dependent upon the cable length.

Essentially, handover occurs on a per-flow basis, so per-flow call admission for the reservation needs to take place again, even though once the call admission is complete, the data are treated on an aggregate basis with other traffic in the same class.

8.4 Proposed Internet QoS Mechanisms

This section briefly examines the QoS mechanisms that have recently been developed within the IETF. Although these are well advanced, and indeed some implementations have been made, they are still open to development. As usual, attention will be focused on the wireless and mobile implications.

8.4.1 IntServ

The Integrated Services (IntServ) solution provides hard guarantees on the QoS experienced by individual traffic flows. It does this through the use of end-to-end signaling and resource reservation throughout the network. The reservation is regularly refreshed. The IntServ architecture provides three basic levels of service:

- The Guaranteed Service gives hard QoS guarantees with quantified delay and jitter bounds for the traffic. It also guarantees that there will be no packet loss from data buffers, thus ensuring near-lossless transmission. This Service is intended to support real-time traffic.
- The Controlled Load Service makes the network appear to be a lightly loaded besteffort network. This class is aimed at delay-tolerant applications.
- Best Effort (no reservation required).

To achieve this, IntServ requires per-flow state to be maintained throughout the network.

It was developed with the assumption that QoS was primarily required for large-scale multicast conferencing applications. This led to the decision to use delay-based admission control. The best-known problem with the IntServ approach is its poor scalability, because per-flow state needs to be maintained in the core network, where thousands of flows may exist simultaneously. Also well known is that reservations need to be regularly refreshed, which consumes valuable resources, especially in bandwidth-scarce environments. Other problems are that the call admission procedures and the routing scheduling schemes are complex and rely heavily on the per-flow state. This is primarily a result of the use of delay-based admission. For example, most IntServ buffers use weighted fair queuing management schemes. The guaranteed service may also lead to inefficient network use, especially within the core network. Whilst many Internet Drafts have been published trying to overcome these problems, the now disbanded IntServ working group issued a position statement saying that the approach is only suitable for small networks. A further question that has been raised with IntServ is that it provides absolute guarantees and essentially duplicates some of the functionality already provided in RTP such as jitter control. Thus, an application that used RTP to provide stream synchronisation would receive jitter control twice if it also wanted packet delay controlled. This is another indication that it is overly complex. TCP is not the only transport layer service that assumes near-lossless transmission. This assumption also underlies the real-time QoS class definitions. For example, the IntServ guaranteed service assumes that, if a router buffer is not overfilled, the delay can be known and all loss avoided. This is not sufficient in the wireless environment, where there is an (uncontrollable) relationship between transmission delay and loss.

8.4.2 Multi-Protocol Label Switching (MPLS)

MPLS was originally presented as a way of improving the forwarding speed of routers, but it has capabilities that enable an operator to provide service differentiation. It is becoming widely used as a network management or traffic engineering mechanism. It appears particularly suited to carrying IP traffic over fast ATM networks.

The basic principle of MPLS is that routers at the edge of the MPLS domain mark all packets with a fixed-length label that acts as shorthand for the information contained in the IP packet header. This label identifies both the route that the packet needs to take through the MPLS network and the quality of service category of the packet. MPLS packets follow predetermined paths according to traffic engineering and specified QoS levels. The label is very short (32 bytes). Thus, once within the network, packets can be routed very quickly on the basis solely of the label. This requires significantly less processing than routing based on analysis of an IP packet header.

MPLS can be used to provide a wide range of different service classes, which could include reliable data transport and delay-sensitive transport services. This service is guaranteed not on an end-to-end basis, but only across the particular MPLS domain. This is a prioritization service, and service level agreements are typically used for admission control. MPLS has received much favourable press, and is used successfully in certain circumstances to provide some QoS today. Equally, however, it has received unfavourable press. Concerns include the amount of processing that is required to turn IP packets into MPLS packets, scalability, and the impact on routing protocol, and finally

security. The security concerns primarily apply to the use of MPLS to provide Virtual Private Networks (VPN), and they are twofold. First, MPLS for VPNs does not, by default, encrypt everything, it is up to human operators to configure the system correctly - and most security problems occur as a result of human error. Second, the humans responsible for the configuration are typically the ISP, not the actual person/ company using the network. This highlights the key issue for MPLS it essentially moves intelligence back into the control of the network operator, breaking away from the end-to-end principle.

In many respects, MPLS for QoS is similar to the DiffServ approach presented below, although, to reduce the scalability problems, it is usually used as a Layer 2 rather than a Layer 3 solution. It provides improved granularity of service at the expense of more complex administration. In itself, it cannot provide end-to-end QoS configurable on a flow-by-flow basis.

8.4.3 DiffServ

The Differentiated Services (DiffServ) architecture aims to provide service differentiation within the backbone networks. It provides a simple QoS, with no signaling mechanism and QoS delivered only to aggregated traffic classes rather than specific flows. Essentially, on entry to a network, packets are placed into a broad service group by the classifier (Figure 8.11), which reads the DiffServ CodePoint (DSCP) in the IP packet header (Figure 8.13), the source and destination address, and determine the correct service group. The correct group or class is determined through static service level agreements (for example, packets from the boss are always given the highest priority). The packets are then given a suitable marking this may involve changing the DSCP. Traffic shaping may then occur, for example to prevent large clumps of data with the same DSCP entering the network. All packets within a specific group receive the same scheduling behaviour. These behaviours can be simple to implement using class-based queuing. Once within the network, routers only have to forward the packets according to these network defined scheduling behaviours, as identified through the DSCP. The complex processing (classification, marking, policing to ensure that no class is oversubscribed and traffic shaping) only takes place at the boundaries of each network domain. This may be done individually by the traffic sources, edge nodes, or a centralized bandwidth broker may be involved. This is sufficient to protect the network and guarantee the service for the aggregate class.



Figure 8.11: Components in a DiffServ border router.

A number of different classes have been defined. These include the Expedited Forwarding (EF) class, which aims to provide a low-jitter, low-delay service for traffic. The definition of this class is currently being tightened, primarily to ensure that it can be easily used within the Integrated Services over Specific Link layers (ISSLL) framework described below. Users must operate at a known peak rate, and packets will be discarded if users exceed their peak rate. The Assured Forwarding (AF) classes are intended for

delay-tolerant applications. Here, the guarantees simply imply that the higher QoS classes will give a better performance (faster delivery, lower loss probability) than the lower classes. These classes have cross-Internet definitions. Finally, network operators are also at liberty to define their own per-hop behaviours - use of these behaviours requires packet re-marking at network boundaries.

This appears to be a good solution to part of the QoS problem as it removes the per-flow state and scheduling that leads to scalability problems within the IntServ architecture. However, it provides only a static QoS configuration, typically through service level agreements, as there is no signaling for the negotiation of QoS. As with MPLS, end-to-end QoS cannot be guaranteed. DiffServ was only ever intended to be a scalable part of an end-to-end QoS solution.

8.4.4 ISSLL

The Integrated Services over Specific Link Layers (ISSLL) working group was initially formed to consider how to provide IntServ over specific link technologies, such as a shared Ethernet cable. One of the key ideas to come from this working group is an approach to provide IntServ QoS by using DiffServ network segments. This solution maintains the IntServ signaling, delay-based admission and the IntServ service definitions. The 'edges' of the network consist of pure IntServ regions. However, the core of the network is a DiffServ region, and all flows are mapped into one of a few DiffServ classes at the boundary depending upon the implementation, in either the edge or border routers of [Figure 8.12](#).




Figure 8.12: ISSLL architecture.

This approach essentially treats the core of the network as a single (logical) IntServ link. This 'link' is created by tunnelling (or IPv6 source routing) the data and signaling messages across the DiffServ Core. This ensures that routing table updates in the core do not lead to changes in the border/edge routers used by traffic. Traffic conditioning may exist both at the edge of the network and at the DiffServ network boundaries.

The advantage of this solution is that it allows hop-by-hop call admission, and flow-based scheduling at the edges of the network, where low traffic densities make this the most practical way to achieve good-quality guarantees. In the core of the network, the scalable solution of DiffServ scheduling can be used, where hard guarantees on QoS can be made on the basis of more probabilistic judgements. From the above discussion, it can be seen that most QoS architectural solutions may be based around the ISSLL solution, with attention paid to the class definitions. This flexible approach, only standardised in late 2000, should finally enable end-to-end QoS for the Internet. Already, small

RSVP/IntServ networks exist, whilst larger network operators are implementing DiffServ core networks.

8.4.5 RSVP

The Resource Reservation Protocol, RSVP, is a mechanism for signaling QoS across the network. It is a key element of both IntServ and ISSLL approaches described above. Although it is strongly associated with the IntServ architecture, it is a more general QoS signaling protocol. Whilst not widely interpreted by routers within networks, RSVP has been widely implemented on a range of different terminals, including Microsoft Windows.

RSVP is an out-of-band signaling system that operates in a soft-state mode although the protocol is flexible, and it is possible to operate RSVP in a near-hard-state mode across any section of a network. This is a particularly useful feature in wireless networks, where it is important to minimise the amount of signaling to save both wireless network bandwidth and mobile battery power. RSVP messages are sent end to end, but carry a flag to enable them to be read and processed by network elements. RSVP assumes that the receiver is responsible for establishing QoS (and, by implication, paying for the level of QoS it receives). It is a two-pass protocol. This ensures that it can handle asymmetric paths, and it enables the sender to identify the nature of the transmitted traffic to the receiver (so that the receiver can make an informed choice as to the level of QoS requested). Whilst the receiver is typically responsible for the actual reservation, the sender also implicitly acknowledges the suitability of the reservation and so can be held responsible in any case of dispute. The sender initially describes the data and identifies the data path. The receiver then sends a reservation message back along the data path - to achieve this, state is installed throughout the network. Initially, RSVP was designed to operate on a hop-by-hop basis, but the ISSLL community has now considered the use of RSVP across DiffServ domains, where only the edge nodes interpret the RSVP messages.

Details of RSVP Signalling

The main messages are PATH and RESV, which establish a reservation, and PATH_TEAR and RESV_TEAR, which delete a reservation. The PATH message is generated by the sender and propagates through the network to the receiver, gathering information about the network. Each node that processes the message records the session identifier and the address of the previous (RSVP enabled) router. The RESV message, sent by the receiver, actually chooses the required QoS and establishes the reservation. This message is propagated back along the same path though the network via each of the previous router addresses, as stored during the PATH stage. This is needed because, typically, within the Internet, messages that flow in opposite directions between two terminals will follow different paths. Thus, without this support for reverse routing, control messages from the receiver would not reach the nodes in the data path. PATH_TEAR and RESV_TEAR messages delete the reservation the PATH_TEAR message is generated by the sender, whereas the RESV_TEAR message is generated by the receiver. This also is propagated using the previous router address mechanism. The process is indicated in [Figure 5.13](#).

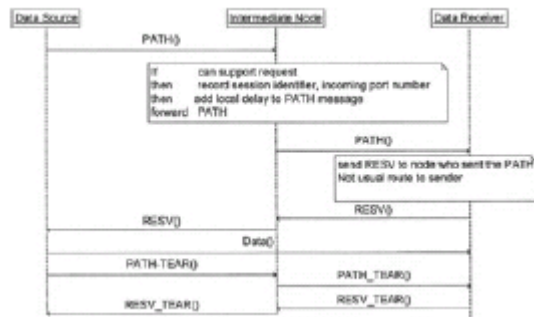


Figure 8.13: Establishing a uni-directional RSVP reservation.

The message contents consist of a number of objects including the session identifier, and the previous (RSVP enabled) hop address. However, most of the message objects, in particular the following, are defined not by RSVP but by the IntServ standards:

- The sender's description of the traffic characteristics (TSPEC).
- The receiver's desired QoS (Flow Spec).
- The network's description of the capability of the Path (ADSPEC).

The traffic description (TSPEC) is supplied by the sender and carried in the PATH message. Since RSVP was developed specifically for multicast applications, this TSPEC is not altered, even if the protocol discovers a network bottleneck.

As the PATH message propagates through the network, the network builds the ADSPEC objects. There is an ADSPEC object for each service that the network supports, and it indicates the amount of resources that are available for that type of service. It is up to the receiver to determine the best service for its requirements. The network information contained in the ADSPEC includes:

- If there is a non-IntServ hop.
- The maximum transmission unit (MTU) size.
- The minimum path latency - Zero if no information is available. This is a representation of the expected (distance related) transmission delay.
- The path bandwidth estimate - The amount of bandwidth the receiver could ask for within the service; this bears no relationship to the bandwidth the sender might want.
- Parameters that enable the receiver to calculate routing delay.

The receiver uses the PATH information to determine what reservation it should make in a RESV message. This is described in the FlowSpec object. This message must be forwarded to each router in the data path using the reverse route installed during the PATH set up stage. Within the ISSLL architecture, the DiffServ region must append a DiffServ class object to the message that tells the sender (or previous node) which DiffServ class to use.

Use of RSVP in a Mobile Environment

In the section on QoS management after handover, a need was identified for a process that repairs the reservation after handover, whilst minimising the signaling and processing load on both the network and mobile terminal. Where the mobility is managed through manipulation of the routing tables (as in the per-host forwarding mobility

management schemes), the RSVP local path repair mechanism is an example of a suitable process. As in the case of handover markings, this assumes that the divergence/convergence nodes hold per-flow state. When a node detects a change in the set of outgoing interfaces for a destination, RSVP can update the path state and send PATH refresh messages for all sessions to that destination. The delay between detecting a path change and sending a path change message is configurable and should be adjusted to give the mobility management mechanisms a chance to build the path. Once the new path message reaches a node that recognises that the message is a result of local path change, it should send a RESV message immediately - thus, the end nodes need not know that the path has changed. Essentially, local path repair is using the detection of a routing change rather than a timer to initiate the soft state refresh messages. It enables quick re-establishment of QoS.

There are a couple of potential problems with using this in the mobile environment. The first is that the mobile node is always either the divergence or convergence node or so using straight RSVP local path repair; this mobile node would have signaling and processing requirements placed upon it. Where the context transfer protocol is used, this situation can be avoided (Figure 8.14).




Figure 8.14: Context Transfer Protocol and RSVP.

Figure 8.14 also indicates that the old reservation is explicitly removed in this process. This is not part of the standard RSVP implementation, which relies on unnecessary reservations being removed through the soft-state management. However, in bandwidth-restricted networks (mobility and wireless networks), this process may not be sufficient. RSVP may be operated in near-hard-state mode to minimise the amount of signaling that is needed within the network. This could then result in 'hanging reservations' being left after a mobility event.

Such hanging reservations could also be left if a session is incorrectly terminated. Thus, if RSVP is used in near-hard state mode through the network, additional mechanisms need to be in place to protect the network. An example of how to achieve this could be to use data traffic as an implicit reservation refresh indicator. The approach described above essentially fixes the reservation after the handover has taken place which leaves the problem of what happens to data whilst the reservation is being repaired. Other approaches to the handover problem in RSVP have also been devised, essentially using the active and passive reservation approach previously outlined. These ensure that a reservation is in place as soon as handover occurs - although with the penalty of additional complexity and scalability problems.

8.4.6 Summary

This section has looked at some of the mechanisms that have been proposed to enable the Internet to provide real-time QoS services. After reading this section, the reader may feel that there is no complete solution for this problem that will work in a fixed, let alone mobile, environment. IntServ is too complex and non-scalable, MPLS and DiffServ do not provide full end-to-end QoS solutions, and ISSLL is a framework for a solution, not a solution in itself, and relies on RSVP, which in turn needs some fixes to work well in a mobile environment. However, many of the required elements are present, and solutions exist for many of the potential problems. The following section therefore looks at one way in which a solution can be created. Within all this, however, it is worth remembering that both DiffServ forwarding behaviours and RSVP are being modified within the IETF, to reflect how people actually use them.

Conceptually, each individual flow is in its own queue, and a certain discipline determines the order in which queues are served and the amount of service time that they have, so that, for example, flows that have requested larger bandwidths will obtain a longer service time. Where thousands of flows exist, this can become complex. Because each flow is in its own queue, this ensures that no flow is adversely affected by the behaviour of other flows.

Here, each aggregate obtains its own queue, and then some discipline determines the service of the queues. This discipline is fixed, and not changed every time a new flow arrives. The queues themselves are served in a traditional 'first in, first out' format, so traffic flows within the same aggregate can affect each other.

Services here are typically the IntServ guaranteed and controlled load services.

8.5 IP QoS for 3G - A Possible Solution

This section now describes one way in which a network operator offering wireless access and mobility support to its users could extend their domain to support QoS. This idea is one of those generated within the EU BRAIN project. This is certainly not the only approach that could be taken, nor is it necessarily the correct approach - after all, many of these ideas are still being developed in the research arena.

The main focus of this network QoS mechanism is to provide one, real-time, service in addition to the normal best effort service. Other DiffServ-based prioritisation QoSs can be easily added to this domain and managed through service-level agreements that exist between the user, the mobility domain, and other Internet domains. The introduction of such services would not affect the behaviour of the described real-time service. This real-time service requires that data be transmitted across the entire network in less than 200 ms, and that no losses due to network congestion should occur.

8.5.1 Overall Architecture

The overall architecture ([Figure 8.15](#)) is based upon the ISSLL architecture. This is because it is likely that ISSLL will be the QoS direction taken by the Internet community already, terminals and local area networks exist that implement IntServ, and core network

operators are implementing DiffServ-based core networks. In keeping with this, RSVP is used as the signaling protocol for real-time services.




Figure 8.15: Architecture for QoS in mobile network.

Within the mobility operator's domain, it is assumed that a per-host forwarding mobility management scheme is used. Such a scheme minimises the use of tunnels. This mobility management zone is assumed to be the same as the region of the network where a per-flow QoS state needs to exist at every router. Since this is the edge of the network, scalability issues are minimised. This is not the only option but simplifies the following descriptions.

The tables then summarise the design choices made, [Table 8.1](#) covering the generic QoS choices and [Table 8.2](#) focusing on specific mobility and wireless issues.

Table 6.1: Summary of generic design decisions	
Topic	Choice Made
Layers of Quality - Transport Layer	All parameters except delay can be managed at this level. Good options are RTP for jitter and stream synchronisation and TCP for loss
Layers of Quality - Network Layer	Delay is the only critical parameter to manage at this level
Layers of	Delay control to support the network layer QoS. Additional loss control

Table 6.1: Summary of generic design decisions	
Topic	Choice Made
Quality - Link Layer	mechanisms assumed
Routing	Aggregate, DiffServ scheduling for speed and scalability Scheduling complexity is minimised, with simple FIFO queues by using bounded delay service for hard real-time traffic
Call Admission	Hop by hop admission for the reservation based traffic in mobility zone. This enables firm QoS guarantees to be made. Edge (gateway) router admission control assumed for backbone, where statistical effects can be used to strengthen QoS guarantees Peak bandwidth only required, but full IntServ parameters carried for interoperability
Signalling System	Out of band to save bandwidth Dynamic (ie per flow) RSVP for reservation based service. Service level agreements for prioritisation services Hard state signalling may be used on a hop by hop basis
Table 6.2: Shows mobility and wireless design choices	
Topic	Discussion
Coupling between QoS and mobility	Minimised. Larger coupling is required to repair reservations in tunnel based mobility management schemes
Behaviour during handover	Context transfer protocol establishes reservation quickly over wireless link. A special DiffServ class is used by traffic during handover to enable it to access a static guard band
Behaviour after handover	RSVP local path repair, as described above is used to restore reservations
Wireless Considerations	Considered within limitations of layered model, and compatibility with existing protocols Mechanisms internal to routers are the responsibility of wireless router manufacturers

The key to this design is how the absolute guarantee of delay is made whilst only using a DiffServ network. This is achieved by using a special DiffServ class definition the bounded delay service. This service allows the making of hard guarantees on the maximum delay that a packet will experience at any one node, and is discussed in more detail later. The total delay is made up of a number of elements: transmission delay, router delay, and the delay across the wireless interface. At the network layer, the individual maximum router delay for the bounded delay service is easily configurable, and the total router delay then depends upon the number of hops in the system. It is possible to build the network layer QoS using the bounded delay service that enables most of the delay budget to be used for transmission and the wireless link. This does not require that all network domains use the same mechanisms for supporting realtime services.

8.5.2 Bounded Delay Differentiated Service

One of the key differences between this solution and standard ISSLL IntServ over DiffServ is that DiffServ routers are used in the domain at the edge. This provides for easier mobility management by using DiffServ to mark, and preferentially treat, packets belonging to nodes that are in the process of handover, avoiding the need for complex pre-reservation schemes. Also, DiffServ requires simpler scheduling and admission control mechanisms than traditional IntServ. Nevertheless, hard, rather than statistical, QoS delay guarantees can still be given. The bounded delay (BD) service has been proposed as a means to provide scalable, guaranteed real-time data transport within the Internet. It allows flows to have a guaranteed bandwidth and low, quantifiable queuing delay, whilst routing is simply based on traffic aggregates that are identified through the DSCP marking. In the general case, it does not require any per-flow state to be held at routers, and admission control is based on a bandwidth sum, making this much more scalable than the IntServ guaranteed service, for example.

Basic Operation of Bounded Delay Service

For each output port, a node has a certain amount of bandwidth that is allocated to this service. Provided this bandwidth limit is not exceeded (closed queue with the maximum arrival rate of traffic known and less than the departure rate), all traffic using this service at that node has the same, guaranteed, worst-case routing delay. All traffic for this service can be scheduled using simple FIFO queuing algorithms. This worst-case delay is fixed for that output port, and is described by [Equation 6.1](#).

$$(6.1) \text{Delay}_{\text{worst_case}} = \frac{N * MTU_{BD} + MTU_{BE}}{R}$$

Where N is the number of active BD flows destined for the output port, and each arrives on a separate input port, MTU_{BD} and MTU_{BE} are the Maximum Transmission Units of the bounded delay and best effort flows respectively, and R is the link speed of the output port.

In addition to the worst-case delay bounds, the authors propose the use of additional statistical delay bounds, to be used in the core of the network. If a worst-case delay is always used in call admission decisions, this can lead to an inefficient backbone network with calls being refused unnecessarily. This is because the worst-case delay will be very rarely experienced at each stage across the whole network path, especially where there are large numbers of flows, so that it becomes very unlikely that BD packets will arrive simultaneously on every input port. Thus, regions can be defined where different admission control strategies may be used, giving significant efficiency gains within backbone networks. In many cases, a simple bandwidth sum is sufficient to determine admission control, whereas in the worst case, a bandwidth sum per input port is sufficient.

Users of this service must request some of this service's bandwidth. This request is propagated through the network, and resources are reserved at each node. The user then marks the traffic with the required code point and must constrain their traffic to the agreed peak rate. To minimise the peak bandwidth required, a token bucket traffic shaper with the bucket depth equal to the maximum packet size may be used. In common with

other DiffServ networks, traffic needs to be monitored ('policed') at entry to the network. However, other functions such as traffic shaping and marking are not necessarily required.

Building a Network Behaviour from the Bounded Delay DS

The above section has just described how to build a router that can guarantee the maximum delay that a packet will experience at that router. In DiffServ terminology, this is known as a per-hop behaviour. However, users are not so much interested in per-hop behaviours as in the whole behaviour of the packet across the network.

The delay that a packet experiences through the network is the sum of the router delays and the transmission latency. To build a real-time service, the end-to-end transmission delay budget is 200 ms. For a transpacific transmission, the transmission latency will not be less than 80 ms. The use of a wireless network can increase this transmission latency. Wireless networks need time to overcome the very high losses associated with transmission over wireless interfaces. The wireless transmitter typically needs 10–100 ms to achieve wireless transmission this figure depends upon the type of wireless system used and the probability of successful transmission. A period of 20 ms will be assumed as typical for this example. Furthermore, both end terminals could have wireless interfaces. Thus, this leaves no more than, say 80 ms available for total router delays through the network. Internet packets have a maximum number of routers - usually 30 - that they can be transmitted through before the packet is destroyed as undeliverable. This prevents circular routing problems. Thus, the delay budget for router delays should be imagined as being divided between 30 routers. However, the delay budget should not be divided evenly between all routers, as the worst-case delay of a bounded delay router will be higher where the output port is a low-bandwidth link. Otherwise, analysis of the equation above shows that either the maximum packet size (MTU) or number of bounded delay flows simultaneously supportable would need to be severely restricted this is illustrated in [Figure 8.16](#). This figure shows that to achieve a worst-case delay fixed at 2 ms, for 1500-byte sized packets, no BD flows would be possible until the output port had a 50 Mbit/s bandwidth. Note that since this is a result of blocking on the output line caused by a packet already in transit, this is not specific to BD, and can only be avoided if packet transmissions are aborted to allow fast traffic to queue jump.




Figure 8.16: Minimum bounded delay of a node is determined by size.

This service only provides delay control there is no separate jitter control. The maximum Jitter is equal to the maximum (routing) delay. This has implications for the size of buffers that are required to ensure smooth playback of data. It is also worth noting that large buffer storage times add to the overall delay that data experience. Since RSVP

enables the routing and transmission delay information to be gathered separately in the ADSPEC object, large transmission delays do not necessarily lead to large jitters. We can make a rough estimate of the required size of a buffer for real-time voice and video, as illustrated in [Table 6.3](#).

Table 6.3: Buffer sizes required if jitter is not controlled independently from delay		
	Voice	Video
Bit rate	64 kbit/s	2 Mbit/s
Total routing delay = maximum jitter	60 ms	60 ms
Required Buffer size	0.5 kbytes	15 kbytes

Thus, buffer sizes are well within a tolerable range (cf. for example the usual 8-kbyte TCP buffer). However, it is worth remembering that wireless networks may add additional jitter to a real-time stream (although ARQ loss management, a large source of jitter in wireless networks, is probably not suitable for real-time traffic).

8.5.3 Mobility Management

The original BD service description only requires a bandwidth sum to be maintained at each bounded delay node. It assumes that some additional policing is used to prevent denial of service attacks. It further assumes that RSVP is not used - it assumes sender-initiated QoS. This avoids the requirements for per-flow state in every node and ensures that QoS control messages follow the same route as the data. This eliminates scalability concerns and allows this service to be used throughout a core network to provide hard real-time QoS. However, as previously described, per-flow state must be maintained at each node of the mobility domain in order to handle mobility events -each node must police each flow to enable it to identify correctly marked BD traffic that has been re-routed away from the original reservation. Thus, RSVP signaling is no extra overhead. BD is still considerably less complex than true IntServ routers, where more complex scheduler techniques and more complex admission control decisions would be needed.

If a tunnel-based mobility management scheme is used, the simplicity of the QoS mechanism makes it easier to manage QoS in the tunnels. BD does not guarantee flow isolation: flows are treated as aggregate flows. Thus, a tunnel can be created whose total bandwidth is simply the sum of the bandwidth of all the BD flows to that destination. This minimises many of the trade-off complexities associated with QoS in general in tunnels.

8.5.4 Signalling

Reservation-based services are provided through interpretation of end-to-end RSVP signaling messages carrying IntServ objects for traffic description. Although most of the information in the IntServ objects is irrelevant for the service described, this choice is fundamental to building a system that is naturally compatible with end-to-end Internet QoS. Readers may be concerned at the use of RSVP within a mobile network. However, it is worth observing that many of the problems claimed about RSVP are actually problems that relate to how RSVP is used in IntServ networks. For example, RSVP is

scalable, but its use hop by hop throughout a network with regular refresh messages as described in pure IntServ is not scalable. RSVP does not add a huge overhead to traffic. If used in a hard-state fashion, it adds barely a 1% overhead to *heavily* compressed voice traffic.

This section verifies that RSVP messages, as used in ISSLL, will provide a suitable signaling mechanism for the bounded delay service, as described above. One issue is that the standard RSVP traffic descriptor, or TSPEC, provides no indication to the network or receiver that the traffic described is delay-sensitive. To achieve this, the sending node would have to prepare an ADSPEC object only for the real-time (guaranteed) service. The advantage of this is that it enables a mobile sending node to include an estimate of the transmission delay associated with its local wireless link. A further advantage of this is that it will limit the overall size of the RSVP message, as; otherwise, ADSPEC objects about every available network service would be added to the message.

One problem with using standard IntServ RSVP is that, whilst the TSPEC includes a field for the peak data rate, it is allowable to set this to infinity, so this field cannot be relied upon to do call admission. Thus, for the bounded delay service, some heuristic may need to be used, which relates the mean to peak bandwidth. It is fairly easy to ensure that the delay factor in the ADSPEC object is correctly entered. Each node providing the guaranteed service must obey [Equation 6.2](#). The bounded delay node sets the D parameter to represent the fixed worst-case delay of the node. It does not need to add anything to the C value.

$$(6.2) \text{ Delay}_{\text{maximum}} < \frac{\text{Size of bucket}}{\text{Bandwidth}_{\text{requested}}} + \frac{C}{\text{Bandwidth}_{\text{requested}}} + D$$

Where C is the bandwidth dependent delay (in bytes) and D is the bandwidth independent delay (in microseconds).

The main benefit of the ADSPEC object is that it allows a receiver to determine how the end-to-end behavior can be understood from the individual per-hop behaviours that are described by DiffServ classes and reserved, in the case of BD, at each node. A bounded delay node may also set the maximum available bandwidth to the (estimate) of the peak bandwidth (from the TSPEC). This prevents the receiver from asking for more bandwidth than is required in an attempt to reduce the delay. This results from the tension between the IntServ (delay-based admission) and bounded delay (bandwidth-based admission) admission schemes. In an IntServ network, a receiver can minimise delay by requesting a larger share of the bandwidth (to the extent that if a user has access to the entire bandwidth on a link, the routing delay simply depends upon how much data they push down the link). This relationship does not hold in a BD network. Such a relationship is also not good to publicise in general wireless networks, where bandwidth efficiency is key.

8.5.5 Discussions

This section has attempted to identify how QoS might be provided in a network that has mobility support and wireless access. The QoS solution finally proposed integrates easily with the ISSLL framework, ensuring that true end to end QoS should be achievable as networks are slowly upgraded to support QoS. The system can provide both reservation and reservationless QoS. Mechanisms can be included to improve QoS during handover, and it can be seen that manufacturers of wireless network equipment have much scope to provide link layers suited to the transmission of IP traffic.

A fundamental difference between this design and that of current mobile systems is that it assumes that the data receiver is responsible for requesting, and paying for, the QoS provided - because the data received are of value to them. Mobile networks assume that the mobile node will establish QoS in both receiving and sending directions, and that the mobile user will be responsible for any charges that result. Whilst billing has not been explicitly covered, one observation is that, if RSVP is used as described above, the sender is responsible for a traffic descriptor and the first portion of the ADSPEC, whilst the receiver is also involved in the negotiation. One of the parties should know that they are wireless. Thus, it should be possible to split accounting, such that a receiver typically pays, but a mobile user may pay a premium.

Indeed, the actual model for RSVP is 'receiver pays, but sender is ultimately responsible', in the hope that this would prevent junk traffic. In a general mobile network, there is also much scope for application layer proxies (SIP servers, mail servers, etc.) to provide this level of protection. One of the main differences between this discussion and current mobile QoS systems is that the emphasis has been on how end-to-end QoS, including end-to-end reservation-based QoS, may be achieved. Most mobile networks are designed simply to provide QoS within the mobile operator domain. This is in part because the natural QoS on mobile networks is often very poor compared with fixed networks, so this investment brings benefits to users, even if they are not guaranteed true end-to-end QoS. Also, this is in part a reflection of the slowness of the development of Internet QoS. However, Internet QoS is now at a sufficiently advanced stage that many good guesses can be made as to its likely nature. Thus, it is more sensible to introduce a system that can be used end to end, or within the mobility domain, as required.

Within the solution described, it is possible to provide QoS only in the mobile domain by use of network proxies for the end-to-end messages. Current RSVP proxy Internet drafts have had many weaknesses, but this is one area that has been studied under the BRAIN EU project, to which the interested reader is referred for fuller details of localised RSVP. Correct implementation of such mechanisms, which would require some additions to the RSVP specification, would enable them to be used alongside end-to-end QoS, without requiring different protocol implementations or extra signaling overhead.

None of the QoS solutions considered have addressed the soft handover problem ([Chapter 2](#)) of CDMA networks. As a reminder, to manage efficiently handover in CDMA systems, frames are duplicated and sent to the mobile via several base stations. These frames must arrive at the node within about 50 ms of each other. Although the IntServ solution to QoS can explicitly control both jitter and delay, and can handle multicast well, and so could achieve the required delivery, it does this on an end-to-end

basis at an application's request. An application does not want to have to make different QoS requests dependent upon the type of link layer used, nor does it want to have an end-to-end IntServ guaranteed service provided for some non-interactive file transfer, just to handle the soft handover problem. One way to manage the problem is to devolve this to Layer 2, as in CDMA networks. From an IP perspective, this could be considered as using an extended link layer (Figure 8.17) - the link is space-multiplexed rather than, say, time-multiplexed as for Ethernet.




Figure 8.17: "Extended link layer".

5.6 Conclusions

This chapter has concentrated on developing the mainstream IETF ideas on QoS to enable them to operate in wireless and mobile environments. Although often a key driver in QoS systems, the issues that may be caused by multicast in such a mobile wireless environment have been deliberately avoided.

To date, however, only small test-bed networks have been built to verify some of these ideas. This is particularly true of support for real-time services. One particular outstanding issue for IP over wireless QoS is the poorly understood problem concerning interactions between the wireless link and the network layer QoS mechanisms. QoS is achieved by building increasing levels of functionality into the network. In many ways, it therefore goes against the IP principles. One issue for IP QoS is essentially how to achieve high levels of QoS, such as that required for voice services, whilst maintaining the low level of network functionality and remaining true to IP principles. Critics of IP networks believe that achieving the same level of QoS for voice-over-IP as current telephony will always be more expensive than the telephony networks. Conversely, critics of the telephony network claim that those networks are over-engineered, and that they would rather have significantly worse QoS, at a significantly cheaper price! Despite recent good progress, there is clearly some way to go before these issues are resolved.

Chapter 9

IP for 3G

9.1 Introduction

In this final chapter, it is appropriate to revisit the theme that started the report. [Chapter 1](#) outlined some of the reasons why IP should be introduced into 3G networks; this chapter will explain in greater detail the technicalities of how IP could be introduced. One result will be that a network is developed that is much more faithful to the original 'Martini' vision than current 3G incarnations.

This chapter will begin by applying the IP design principles, plus the QoS, mobility management, security and service creation pieces from the preceding chapters, to sketch out a vision of an 'all-IP' mobile network. Of course this will be open to debate and will reveal the author's own prejudices about the meaning of 'all-IP'. This is a serious point, as the term 'all IP' has come to be used in several ways, some of which do not adhere to the concepts outlined in this report. An IP architecture is in fact quite different from the traditional cellular systems that are defined by the network elements, the interfaces between them, and the protocols that run over those interfaces. The IP approach has very weak interfaces and largely concentrates on protocols - typically one protocol providing a single function - which are developed independently and are not tightly integrated to either each other or a particular underlying network structure. Another point is that there are still many holes that IP technology currently cannot fill - areas where work still needs to be done to replicate some of the functionality of the tightly integrated/proprietary standards of 3G.

Having outlined a vision for this all-IP future, this chapter will detail an imaginary journey of a user of said network, seeing how they are able to access all sorts of multimedia services and be able to select network operators based on price and performance. The economic case for IP in 3G was made in [Chapter 1](#); this chapter will concentrate on the potential user advantages and note the compelling similarity of what an all-IP network offers with the original vision of 3G when it was conceived in the late 1980s.

Next, the chapter examines how UMTS is adapting to the IP pressure and critically examines what the next releases (R4/5), often dubbed 'all-IP' in the sales brochures, have to offer and how they compare with the vision. There are also other initiatives on the IP front- including a move to utilise Wireless LANs and, possibly, integrate them with UMTS, which will be investigated briefly.

9.2 Designing an All-IP Network

9.2.1 Principles

How should an all-IP network be designed? By expanding some of the principles described in [Chapters 1](#) and [2](#):

- Layer transparency - The interface between the layers should be clear-cut, and each layer should offer a well-defined service to the layer above. Also, the service does not open the PDUs (protocol data units); it accepts them as unopened packages and acts only on the information given with the PDU.
- A corollary of layer transparency is that layers should not be broken - Layer 7 (applications) are not talking to Layer 2 (link layers) (except possibly to configure them in a management sense). The layers can then be changed independently (e.g. swapping wireless LAN for Bluetooth) without the whole comms software stack needing to be re-written.
- End to end - The terminals do as much of the work as possible; since they really know what they want, it is more efficient to provide the service end to end. Hence, packets always carry the full destination IP address and not just a label or ATM VC identifier. However, there is a need to avoid terminals requiring car batteries. If the access network can reduce the signalling load, that is probably a good thing.
- A corollary of this point is that the transport network should do just that transport, and nothing else. No call control, no unnecessary functionality and the added functionality (intelligence) such as it is, moves to the edge of the network.
- IP networks should be modular- To allow rapid evolution and exploitation in novel ways as well as incremental roll-out. This will only happen if the components are capable of independent evolution/replacement without the need for a complete upgrade of every layer/component. The interfaces between the components should allow freedom of variation. It all works as long as the new protocol has the correct interfaces and performs the required function (e.g. packet delivery).
- IP networks are designed to allow, if required, the value chain to contain several players. There are very few interfaces; the network simply delivers packets, and services are created at the edge. An example of this is the 'dial IP' architecture some network providers allow the user access to a range of ISPs. ISPs in turn allow the user access to a range of online booksellers, and the user can buy items from the Internet with a range of credit cards.
- Mobile networks must reuse as much as possible the transport, protocols, and applications from the fixed world. Mobile always has been, and probably always will be, a small fraction of the total network traffic. Spectrum in the 0.5–3–GHz range is expensive, using spectrum efficiently is complex, and improvements in spectral efficiency (i.e. the bit rate that can be squeezed into a particular chunk of spectrum) are modest. In fixed networks, the capacity of fibre optics is truly vast; Moore's law operates on the transmitters and receivers - meaning that 40 Gbit/s can now be transmitted for much the same cost as 155 Mbit/s a few years ago.

ADSL and Gigabit Ethernet cost a few pounds and so forth. So, the fixed world drives the low-cost, volume transport market - implying that mobile networks should use standard IP routers, i.e. what is used in the fixed world today and interfaces in standard ways. It also implies that for users and developers to gain maximum economy of scale, the same e-mail client should be used - implying the same TCP socket - implying the same IP transport underneath.

9.2.2 Overall Architecture

The first issue is the scale of the network - where does it start and where does it interface with other networks? Without doubt, the network starts at the base station end of the radio link. There is one, and only one, Layer 1/Layer 2 radio hop, then the IP packets are reconstituted (if they were segmented for the radio hop), and then the addresses on the packets are used for the next hop. There is definitely no ATM, AAL2, MPLS or other network layer switching/routing going on. This is a vital point - the BTS, Node B, or transmitter - depending on the terminology used - is an IP router and routes packets. An extensive non-IP network, even a few 'hops', is breaking the IP architecture principles.

The second issue is that there should be a specific access network to conceal all the mobility management and 'lumpy', edge of network QoS, from the rest of the Internet - as well as hiding issues such as the high error rate over the air, and the fact that radio coverage sometimes disappears and so on. What is meant by lumpy QoS? At the edge of the network, if a 1 Mbit/s video session hands over from a neighbouring cell, it can have a great effect on the local resources. If a 3G cell is offering 2 Mbit/s per cell shared between all the users, 1 Mbit/s is a 50% change of resources, and that might imply that there would have to be drastic reductions in bandwidths of six students browsing the web to fit it in. The whole nature of the real-time/non-real-time traffic leaving the cell may have altered. In the core, however, with highly aggregated traffic, the likely changes in traffic load and type are much smaller, and changes take place over longer time scales (e.g. busy-hour). At the core-facing edge of the access network, there should be a normal Internet gateway (running Border Gateway Protocol (BGP) and perhaps acting as a firewall) - in fact, there should be several of them for resilience, scalability, and shorter paths through the network. The access network operator can provide services such as e-mail and web hosting from within their network, or the user can obtain services from any service provider through the Internet.

Figure 9.1 shows a first attempt at the architecture - instead of Node Bs, there are Mobile Access Routers; the Gateways corresponds roughly to a GGSN.




Figure 9.1: Outline all-IP mobile network architecture.

9.2.3 Routing and Mobility

Clearly, mobility is needed in a network, so that users can be reached for incoming sessions and so that a session (such as voice) can continue when a user hands over across access routers. Following the discussions in [Chapter 7](#), the problem can be broken down into three parts - paging, routing updates, and signalling between the access routers. For the final item, a promising approach is the IETF's 'fast mobile IP' approach of having a temporary tunnel between the access routers. However, there is still a choice of how to do the routing updates.

In this case, a per-host-forwarding scheme is most likely the 'purest' solution, IP architecturally speaking. This is because tunnelled solutions, like Mobile IP and its variants, are an admission of failure. The underlying network does not deliver the packets properly (its only real job remember from the discussion above), and so a new tunnel anchor (e.g. the Home Agent) must be employed. The packets themselves are hidden away, and so are their headers - so things like QoS and security are difficult. Also there is the limitation of future evolution (and present services): how will multicast work if all the traffic to mobile users is being tunnelled? One tunnel will be needed per user, and the multicast (or anycast) advantage is lost. Suppose 90% of future mobile services turn out to require multicast. A per-host scheme does exactly what an IP network should do, according to the IP principles: deliver packets.

How are hosts identified? They will need an IP address belonging to the access network where they have roamed, and the address needs to be globally routable. This address must be given up when the host leaves the domain (i.e. the routers within the ownership of one organisation).

Now, having established that an IP address is needed, it will be received either when the user signs on or only when the user starts transmitting or receiving packets. The address will be returned perhaps when the user leaves the domain, or when they have finished their session (i.e. no more data for a 'while'). It would be unlikely for an address never to be returned (i.e. becoming a user's permanent id), since domain owners will not want users walking off with their addresses. It turns out that per-host schemes work best by limiting users to having a valid IP addresses only when they are in an active session, i.e. they do not have one whilst in idle mode otherwise a lot of state (spaghetti routing) builds up in the routers from tracking terminals that are moving but not communicating. This implies that paging (i.e. alerting an idle mode terminal) cannot be done on IP addresses.

Now for the controversial part .Imagine a session layer id, a SIP URL - sip:dave.wisely@bt.ums. In this scheme, paging is triggered by SIP INVITE messages being forwarded from a user's home domain's proxy server (in this case, the domain bt.ums). When the user enters the IP Access Network (AN), they register this with their home domain SIP registration server - meaning that INVITE and other SIP messages are forwarded to the AN. The SIP proxy in the A consults its local registration server to discover the user's paging area identifier which could be the multicast address for that group of access routers. When the user has been paged, they acquire an IP address and report this back to the registration server, allowing the INVITE message to reach the user.

In addition, there is no need for paging for mobile-initiated sessions - the user just acquires an IP address in the same way as dial-up works today. Some say this is not at all in the spirit of the IP architecture, and that anybody should be able to send IP packets of any type and to receive them. For example, home agents always have a care-of address to send even a single packet. This is rather pointless in mobile networks, as it will always come at a cost to the network performance. If there is no filter for packets being sent to mobile terminals, a user could, quite reasonably, start an instant message service that pings all the registered terminals every 10 min (say). If the terminals have to be idle for 15 min before they stop sending location updates at the cell level and move to updates on paging area - that user is greatly increasing the signalling traffic. Also, that mobile user, may be paying per packet delivered and will object to junk mail and packets arriving and, perhaps, asking for high-quality QoS at vast expense to view a margarine video ad. SIP, with its user contact preferences, is well placed to act as a filter in this situation and to trigger paging.

9.2.4 Quality of Service

For the QoS solution, there are still some very difficult questions, such as:

- How can end-to-end QoS be ensured when the end-to-end path crosses several domains?
- Can any QoS be provided in the absence of end-to-end QoS?
- How can QoS be achieved in the face of the particular problems raised by mobility and by the wireless environment?

The end-to-end QoS problem will be solved for the fixed network and is not a mobile specific issue. It could be fixed tomorrow if everyone agreed on the signalling and service level agreements, and maybe users are simply waiting for a de-facto standard to emerge. However, this process could take a long time. In the absence of end-to-end QoS, the access network (AN) might be expected to be the weakest link, because, for instance:

- Its capacity is restricted.
- The QoS over the wireless link is poor.
- Handovers disrupt any QoS reservations.
- Unpredictable mobility patterns make dimensioning, traffic engineering and admission control harder.

Therefore, it is not unreasonable to suggest that QoS might be required within an AN, in order to enhance the effective overall QoS.

As discussed in [Chapter 7](#), a promising approach to providing QoS in the Access Network is based on the Integrated Services over DiffServ architecture. In order to deliver QoS for realtime applications, the bounded delay service is used, with RSVP signalling to reserve bandwidth at the required routers. In order to deliver QoS in the Access Network when end-to-end QoS is absent, [Chapter 7](#) suggested introducing a proxy, and using a modified version of RSVP called 'localised RSVP'. This allows the mobile terminal to initiate the outbound QoS set-up and to instruct the proxy node to initiate inbound QoS. This allows QoS not only within the Access Network for multimedia services but also for activities like web browsing - where the web server will not pay for QoS, but the mobile might be prepared to. As regards QoS over the wireless link, this must involve co-operation between layers 2 and 3. Later we describe a powerful new interface between the two that would provide some of this co-operation.

9.2.5 Security

For security, there are two important issues: mutual authentication between the terminals and network, and confidentiality.

For mutual authentication, there must be a shared secret between the terminal and whoever is doing the billing. So, in GSM and UMTS, it is the SIM card, for a customer and a bank it is the customer's signature, and so on. In an IP network, it does not matter what the secret is – it could be a 128-bit key buried in a card (this is much safer than in the memory of a computer, say - where hackers have shown that it is easy to overwrite/steal keys and passwords). There must also be a security association between the service provider and the access network provider - so that the service and network providers can exchange keys/challenges and the terminal can then challenge and authenticate the network and know that it is 'approved' by their service provider. In practical terms, this means that there are AAAL (AAA Local) and AAAH (AAA Home) servers that are able to exchange details about the subscribed services for which the user is entitled to be billed (QoS class, credit remaining, and so on). The local AAA server also needs to trust the access routers - since they must accept (and authenticate), probably in real time, handovers from mobiles.

If the IP end-to-end principle is followed, confidentiality should be provided end to end using something like IPSec. This now represents less than a few per cent performance overhead on modern machine - and, in the future, even less (Moore's law). However, there are reasons why many transactions would not use end-to-end security (e.g. due to the processor cost of encryption on small terminals or difficulty of compressing encrypted headers), and so the network should also provide encryption over the air to prevent casual eavesdropping.

9.2.6 Interfaces

There is a need for inter-layer interfaces in a modular IP network - both to allow interoperability and to partition the problem (e.g. confining the mobile-related issues to the RAN). Traditionally, IP service interfaces have not had complex functionality, but enhancing them is a way to preserve layer separation, maintaining the IP principles, whilst enhancing performance to reach the functionality of traditional mobile networks. The most important interface is probably the so-called 'Layer 2.5' - between the air interface and the network (IP) layer. An all-IP network should be capable of connecting to many different air interfaces (e.g. WLAN and TDMA), so a generic Layer 2 to Layer 3 interface is needed. Moreover, it has to have considerable functionality if the overall performance of the system is to be efficient. For example, handover can be done entirely at Layer 3 - using only IP messages. However, it is the network card and Layer 2 processes that measure the signal-to-noise ratio and know that handover is soon required. [Chapter 7](#) showed that such Layer 2 hints can greatly improve the performance of handover (packets lost, packets delayed, etc.). Similarly with QoS, most wireless link layers have buffers with a QoS mechanism, and wireless LAN access points might operate a Call Admission Control process. All of these must work in conjunction with the IP layer processes. For example, there is no point in handing over a call to find that there is no Layer 2 capacity for it or making detailed end-to-end QoS signalling and set up to find that the link layer, across the air interface, cannot support the QoS. There is a trade-off, then, between doing everything at Layer 3, but inefficiently, and doing something with the help of Layer 2 but needing a complicated interface to do it.

One such interface that has been proposed is the IP2W (IP to Wireless) interface developed within the EU BRAIN project (www.ist-brain.org) ([Figure 9.2](#)). Each function has associated primitives that allow it to be used in a generic way, and a convergence layer then adapts each underlying link layer to provide the functionality. A discovery function allows the terminal and access router to find out which of the optional functions are supported (e.g. whether Layer 2 encryption is offered).

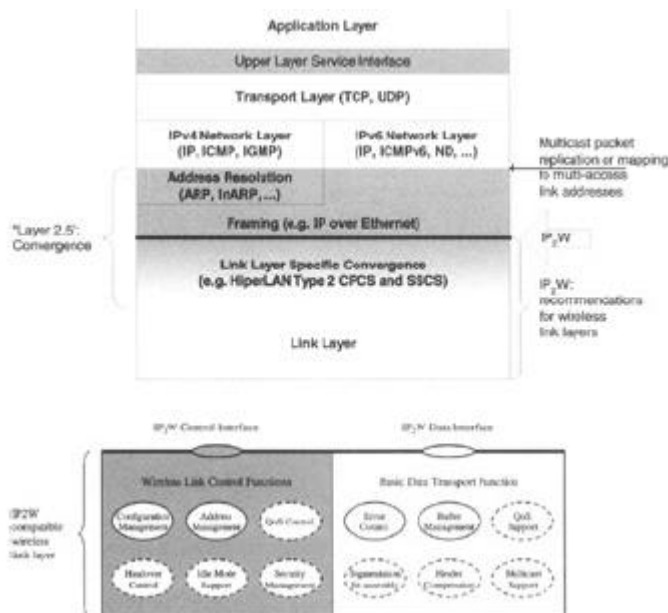


Figure 9.2: The IP2W interface, specified by the EU BRAIN project.

Another interface is clearly the transport service interface offered by the transport layer to the applications. According to one of our IP design principles - keep layer transparency – nothing above the transport layer is allowed to know the details of how the packets are transported. Of course, this is not true today, although one could argue that the socket interfaces are an attempt at producing this functionality, albeit at a very low level. As networks become more complex, better standard simple interfaces will be required. Higher layer components, often referred to as QoS brokers, can then use this functionality for managing the network resources as well as coupling it with local computer resources (e.g. memory or CPU time) to achieve greater QoS. However, all of these issues are above the network design issue focused on here.

9.2.7 An Answer

Figure 9.3 shows an all-IP wireless network. This is an adaptation of a picture that began in Eurescom Project P810 - about replacing ATM with IP in wireless networks. The intelligence is provided at the edge of the network and is split between the access network provider (AAAL, SIP proxy for paging and DHCP for address assignment) and service provider (SIP proxy for service provision, including personal mobility, AAAH for accounting and billing engine and DNS). Of course, there are many missing details, but there is only space here for 4000 words and, after all, 3GPP have used 4000 pages for the complete R3 UMTS standard!

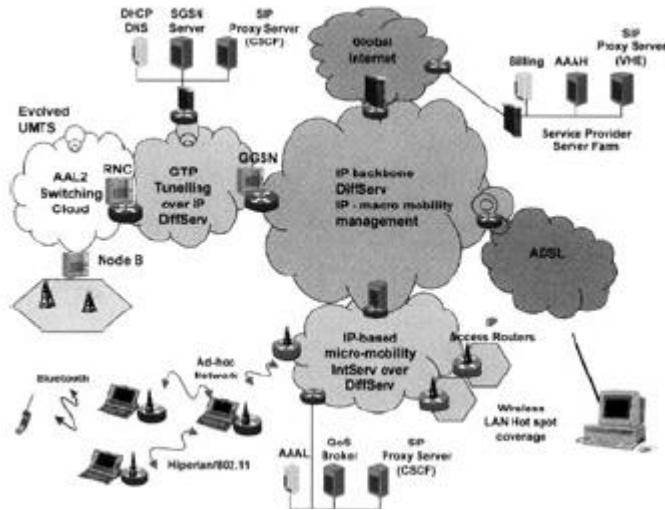


Figure 9.3: An all-IP wireless network.

9.3 Advantages of an All-IP Network

Returning to the imaginary user from the 3G chapter, Mary (only by the time an IP network has been rolled out, she has finished her Ph.D. and is on the teaching staff at the University), this section examines what advantages she might gain from using this all-IP network.

Mary starts her day at the University, where she is contracted to lecture one day a week, by powering up her mini laptop - this is equipped with wireless LAN (WLAN), Bluetooth and GPRS network cards and is set to scan for the available networks. In the University, WLANs are available in many parts of the campus, and so Mary's laptop chooses the University as her access network provider (lowest cost) and presents her SIP URL `sip:mary.jones@x-tel.com` and the AAAL (Local Access Authentication and Accounting server) contacts xtel to authenticate her and the University network to each other. Details of Mary's subscription - Silver Class - are downloaded to AAAL and hence to the access router she has contacted. Mary wishes to check her e-mail and so acquires an IP address locally. An incoming instant message is sent to her by her friend Bob, in the form of a SIP message to her SIP URL - this is redirected from Xtel's SIP server to the University SIP server and, because the DHCP process had created an entry in the University SIP server, delivered to Mary.

Next, Mary wants to start her multicast teaching application - all the students join the group and shared applications run over the top. Because the access network handles multicast properly, the multicast tree is very small - if she had been using UMTS or GPRS, each user would have required a GTP tunnel from the GGSN.

Finally, the lecture ends, and Mary sits in the cafe having a much needed cup of tea. She idly browses the web for Bob's birthday present, typing in URLs from the University magazine, unaware that the pages she is looking at come from a local web cache and not a distant server.

Only because the IP packets themselves are available locally rather than encapsulated - it is this caching possible ensuring a quicker, cheaper service. When she finds something

she likes, she buys it with her credit card - a smart card that fits into the back of her laptop, and that sets up an IPsec connection to her credit card provider. Mary makes a voice over IP call to Bob - her terminal uses RSVP signalling to set up the end-to-end QoS over the variety of networks used by the call. The University access network uses ISSLL (IntServ over Specific Link Layers - in this case IntServ over DiffServ), and the core network uses pure DiffServ. Bob is on a UMTS network - requiring him to set up the QoS for his leg of the connection with PDP context messages. These set up DiffServ markings in the UMTS core network and AAL2 and radio bearers in the Radio Access Network.

Unfortunately, after a series of seamless handovers between different WLAN base stations, Mary wanders out of WLAN coverage and so her terminal executes a vertical handover to UMTS. First, it gains a UMTS IP address, then it sets up a PDP QoS context and uses SIP to INVITE Bob to the same session on the new IP address.

Finally, Mary must attend a meeting of the department staff - the meeting consists of six people in the room and one person who is in another building. Mary's laptop is used to connect, via the WLAN, to the distant colleague - through the University Intranet - and the others all connect to her laptop by forming an ad-hoc network using Bluetooth. Mary's laptop then acts as a mobile router relaying IP packets to and from the Internet - after downloading the appropriate mobile router software at the start of the meeting. After work, Mary goes home, switches off her laptop, and reads her post. How does this all-IP network compare with the original UMTS vision from the late 1980s discussed in Chapter 2? It certainly offers a variety of access technologies including cellular, Wireless LAN, Bluetooth and ADSL. It offers true broadband connectivity with WLANs such as 802.11 and Hiper-lan 2 (10 Mbit/s+) in some hot spots. A SIP-based VHE could also allow common service to be adapted to location and access technology (e.g. bandwidth).

So, in the sense of the user functionality, it probably is closer to the original vision than the early versions of 3G networks. However, it does not include a satellite component and it is not the universal system envisaged. Of course, there are many difficulties and unresolved questions to moving to such a network. In addition to the issues mentioned above, such as paging and protection against spam, there are also other issues such as:

- Whether soft handover (for CDMA systems) can be supported on an all-IP Network without a specialised Layer 2 switching network. As seen in [Chapter 3](#), the nature of soft handover in UMTS requires the user data to be delivered to a set of base stations with very tight control of the timing of the streams (less than 100 ms difference). Current IP protocols do not provide a solution to this problem, and without significant additions as seen in the QoS chapter, IP does not provide tight packet jitter control.
- How the network can evolve from the current standards, in order to exploit existing investments and to support existing terminals.
- Whether there is any benefit for operators in allowing transparent connections to other services - and indeed for breaking apart the value chain that is currently so tightly linked to SIM-based authentication.
- What cost advantages such a network brings or whether it is dominated by the spectrum and air interface costs.

Perhaps a good way to predict if and when such an IP network will be introduced is now to look at how new releases of UMTS, currently being standardised, are moving to incorporate more IP ideas, architectures, and protocols.

9.4 3G Network Evolution

The evolving 3G standards currently being worked on involve two new concepts for traditional telecom networks: voice over IP (VoIP) and IP call/ session signalling for multimedia services. This section will provide a brief overview of both.

9.4.1 UMTS R4 - All IP Transport

The second version of UMTS was originally called Release 2000 - following the year that the standardisation was expected to be complete. However, it was soon realised that the changes being made from the original version (R99) were so large that they would have to be split into two standards that would not be complete before 2002. Consequently, the original version of UMTS was named R3 (since it was the third standards release by 3GPP) and the new UMTS versions called R4 and R5.

UMTS R4 is only concerned with the Core network part of the circuit-switched domain (CS Domain) the UTRAN and packet switched (PS) domain remain the same. R4 takes the Iu-CS interface and allows it to be connected to a media gateway so that the voice traffic can be carried in IP packets a form of voice over IP (VoIP). The general architecture is shown in Figure 9.4.

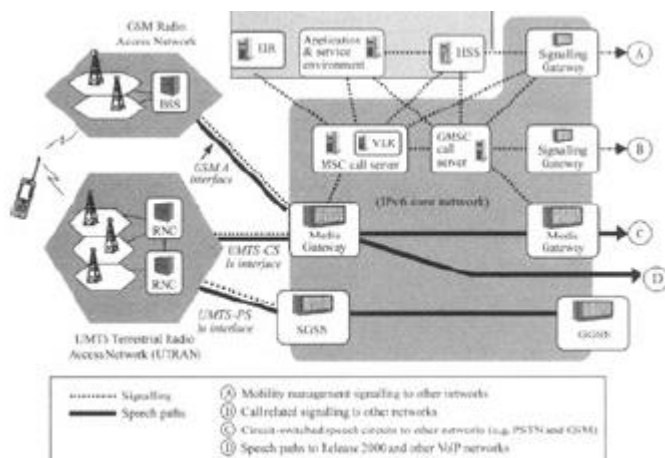


Figure 9.4: UMTS R4 architecture.

The important point about R4 is that it is fully backwards compatible with R3 (R99) the terminals are unchanged and do not require an upgrade; they offer exactly the same services and capabilities. The advantages with this system are the cost savings, integration, flexibility, and evolution. The cost savings are expected to arise from IP providing a cheaper switching technology compared with a core linked by TDM circuits with 64 kbit/s per channel or ATM technology. In addition, in R3, the low-rate mobile speech (Adaptive Multi Rate codecs, giving a variable rate coding from 5 to 12 kbit/s) is converted into 64 kbit/s PCM (Pulse Code Modulation) in the MSC - if this connects to another mobile network, savings could be made by not transcoding the speech (i.e. from

UMTS AMR to 64 kbit/s back to AMR - a major processing overhead). The R4 network offers this flexibility. Cost savings also arise from being able to run both CS and PS domains over the same core, and this increases the flexibility and allows integration of monitoring and control functions, for example. Operators might also have a core IP backbone that can then be used for all fixed and mobile traffic. In R4, it is also possible to dimension the user plane and control plane functions separately - Media Gateways (MG) or Media Gateway Controllers (MGC) can be added independently (see [Chapter 5](#) for a full explanation of gateways and controllers). Finally, R4 represents an evolutionary step towards a full VoIP solution - where voice is packetised in the terminal. It was considered too large a step for operators, manufacturers, and standards bodies to achieve this in a single development.

What has, in effect, happened is that - compared with R99 - the MSC has been split down the middle. The switching and user plane part has been replaced by a MG, and the control, call state, and service logic part has been turned into an MSC server. Signalling from the UTRAN is relayed to the MSC server over TCP-IP - the MSC server controls the MG using the H248/ MEGACO protocol. The GMSC has also been split down the middle - with the GMSC server performing all the call control and HLR interrogation of an R3 MSC. Connection to other networks can avoid the conversion back out of IP packets if the speech paths are compatible.

9.4.2 UMTS R5 - IP Call Control and Signalling

UMTS R5 changes only the packet switched (PS) core network the circuit-switched part of the core can be the MSC/GMSC of R3 or the MSC-server/ MG architecture of R4. R5 introduces two major elements to the PS core network:

- A new core network domain - called the 'Internet Multimedia core network subsystem' or IMS for short.
- An upgrade to the GSNs to support real-time voice and other delay-sensitive services. The UTRAN is also upgraded to support real-time handover of PS traffic but is otherwise unchanged, with the interface between the core network and UTRAN being via the normal AAL5 Iu(PS) interface. The overall R5 architecture is shown in [Figure 9.5](#).




Figure 9.5: UMTS R5 architecture for connection to other R5 networks.

The real purpose of R5 is to enable an operator to offer new services examples might be: multimedia conferences (e.g. voice, video and whiteboard), a multi-player, interactive game, and a location-based service. The IM domain is about services - their access, creation, and payment - but in a way that allows the operator to keep control of the content and revenue.

(There is an interesting contrast between traditional voice networks, where services are integrated within the network and under the control of the network operator, and IP networks, where services are provided at the edge of the network in a way that is decoupled from packet delivery.)

There are three fundamentally new aspects to the IM domain - call/session set-up and control, roaming, and the use of IPv6. The next sections look at each in detail, starting with call/session set-up and control.

Another issue concerns the treatment of voice in an R5 enabled network. Clearly, it could now be carried over the PS domain as VoIP, but that does not mean that it will be. In practice, MSCs (R3) or MSC-servers (R4) will probably carry the voice-only traffic for a long time to come. Amongst the reasons are that: there will be many non-R5 terminals that must be supported anyway; and because the amount of voice traffic is more predictable than for multimedia services, keeping the traffic separate might make network management and dimensioning easier (at least whilst voice is the dominant traffic source).

Call Control

In the R3 packet switched domain, there is no concept of a call or session (here, the term 'call' will be dropped, instead simply referring to multimedia sessions that might include voice calls). Users set up a PDP context and connect to their chosen access point - probably an ISP or corporate LAN. They can access services, such as web browsing and e-mail and even video streaming, but real-time interactive services will not be supported by the offered QoS. As an example of the benefit that the IM domain brings, consider a multimedia conference. Imagine that a user has a laptop and an R5 PCMCIA card, and wants to have a video/voice/whiteboard session with two colleagues something extremely difficult with R3.

They start the session on voice and whiteboard and only add the video half-way through. The IM domain needs to provide the functionality to allow users to: locate each other, share information such as codec type and bandwidth as well as adding new elements to sessions and generating the Call Detail Records (CDRs) for charging. As described in [Chapter 6](#), 3GPP looked at two protocols for session initiation and negotiation - SIP and H323 - and chose SIP largely because it was more Internet-based, and the possibility existed to build on standard SIP to add 3G functionality, developing this extended capability in co-operation with the IETF. IM domain users are identified by a SIP URL - [<sip:mary.jones@xtel.com>](mailto:sip:mary.jones@xtel.com) - this allows an easy way for users to be contacted from IP domains. Note that the IM domain is totally unaware of the mobility of the terminal, since it is outside the packet switched domain. UMTS R5 introduces a new functional element in the network called the call server control function, CSCF. It acts as a SIP proxy server, and its main jobs are to:

- Locate users - translating SIP URLs to IP addresses.
- Proxy INVITE messages.
- Maintain information state on the session - to allow other multimedia streams to be added to an existing session, for charging and because it controls the MRF (Multimedia Resource Function - essentially a bridge that allows multi-party sessions in a network that does not support multicast).

Before any SIP messages can be sent to the CSCF, the terminal must set up a PDP context especially for this purpose (CSCF discovery is covered in the [next section](#) on roaming) this PDP context uses the interactive QoS class and is only used for signalling. The terminal runs a SIP user agent, and the CSCF communicates with it via its IP address and port number. There is also signalling between the CSCF and the GGSN, so that the GGSN can be informed which flows are IM-related, in order to provide them with better QoS than non-IMS flows. The protocol chosen for this signalling is the IETF's COPS (Common Object Policy Service) protocol for policy management. A PDP context appropriate to the multimedia session can then be established. R5 terminals must carry the AMR (Adaptive Multi Rate) speech coder of earlier releases as a default and might, typically, produce a 12.2 kbit/s bit stream.

A remaining problem is that of header compression. Without it, 12 kbit/s of speech becomes 28 kbit/s for example, and the issue is unresolved as to where the compression will be terminated (RNC or SGSN). If, however, the user requires end-to-end encryption, as provided by IPsec, header compression is not possible, and the user would have to pay for 28 kbit/s across the air.

Roaming

In GSM, a user can roam on to visited networks - provided that the visited network can access the home HLR, and an agreement exists between the two operators. The same kind of roaming is supported for R5 multimedia services. In GSM roaming, call control always takes place in the visited network, the only connection to the home network being access to the HLR.

In UMTS, there was a long and complicated discussion about whether IP multimedia call control for roamers should take place in the visited or home network. Those who said it should take place in the home network pushed the argument that the user would have signed up for a range of services, and many of these would not be available or would work differently in a visited network. Those who favoured visited network control were concerned about the long delays and signalling traffic created by having all services controlled from the home network that might be located on a different continent.

In the end, it was decided that R5 control would be controlled from the home network. This complication gives rise to three 'flavours' of CSCF - the P-CSCF (proxy CSCF), S-CSCF (Serving CSCF) and I-CSCF (Interrogating CSCF). To explain the roles of these three CSCFs, suppose a user wishes to make a VoIP (voice over IP) call to a user on another R5 network.

Before using IM domain services, including being able to receive an IM session invite or call, a user must first register with the network. This always takes place via a proxy CSCF (PCSCF), whether users are in their home network or not. The P-CSCF provides

basic multimedia session support as well as acting as a firewall to the IM domain. Users discover the P-CSCF by first activating a PDP context for signalling and registration, gaining an IP address either dynamically or statically in the usual way, and then sending a DNS query for 'P-CSCF' - the DNS server at the GGSN then returns a P-CSCF IP address. All mobile-to-network signalling is sent to a P-CSCF, and the mobile never discovers the address of the other CSCFs.

A REGISTER message is sent by the user to the P-CSCF, and this is relayed to an interrogating CSCF (I-CSCF) in the home network (the home network could be identified by the P-CSCF using the IMSI or SIP URL of the user). The I-CSCF acts as a gateway for foreign networks, polices access to the IM domain via foreign networks, and interrogates the HSS (Home Subscriber Server).

The HSS is simply the HLR with new capabilities added to take into account the IM domain role. The HSS is also access-independent, so that operators can reuse the IM domain for other IP access technologies such as DSL and packet cable.

The I-CSCF in the home network retrieves the data about the subscriber from the HSS, probably using LDAP (this is not yet defined) and selects a CSCF to actually deal with the requested service - called the serving CSCF (S-CSCF).

The serving CSCF actually has much more functionality than the proxy and interrogating CSCFs. It has access to the resources needed to create services, such as video servers and media gateways. An operator may have several S-CSCFs with different capabilities and select the one to handle the session based on the requested service.

The I-CSCF in the home network distributes the data retrieved from the HSS to the CSCFs. At the end of this procedure the P-CSCF knows the IP address of the S-CSCF, and the P-CSCF and S-CSCF have information about the subscriber from the HSS. As in GSM, the HSS is aware of the location of the user to re-route incoming INVITE messages. [Figure 9.6](#) outlines the process for a call using the IM domain between two mobiles attached to different IMSs. The signalling involved has all been detailed in 3GPP standards. (See [Further reading](#) for the details).

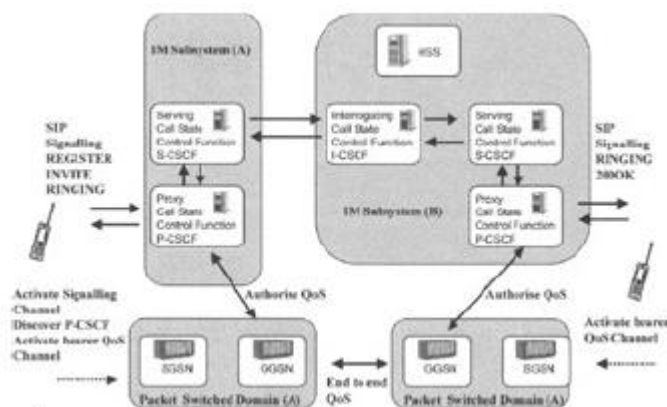


Figure 9.6: Mobile to mobile session flow.

IPv6

The IM domain will use IPv6 - so that user equipment will have to have an IPv6 stack to register for IM services. The fundamental reason for using IPv6 was the exhaustion of IPv4 addresses especially in Europe and Asia. IPv4 addresses are very limited in these areas, and with hundreds of millions of new users expected, it was felt that the IPv6 address space was needed to cope with this. In addition the enhanced security, auto-configuration and header flexibility of IPv6 will simplify the service creation environment ([Chapter 5](#) gives a brief overview of some of the new features in IPv6).

The important issue with using IPv6 in the IM domain, and indeed with IPv6 in general, is the question of interworking with IPv4 networks, which are expected to continue to exist for many years after the launch of R5. When UMTS is launched, it will be based on IPv4 carried over ATM. However, user packets are carried over this network inside the GTP tunneling protocol, and so it is possible to carry both IPv4 and IPv6 packets across an IPv4 GPRS network - provided the GGSN operates a dual stack (i.e. runs both IPv4 and IPv6 stacks).

Hence, upgrading UMTS to IPv6 - to replace the ATM transport and IPv4 functionality, for the sake of a unified core network - could then be undertaken anytime with no implications for users or services. The IM domain would then be reachable from R5 terminals using IPv6 signalling, but it is likely these that would feature a dual (IPv4/IPv6) stack. When interacting with an IPv4 network, e.g. an ISP or corporate LAN, the R5 terminal might choose IPv4 to avoid the need for interworking. Interworking, however, will definitely be needed for the IM domain to interact with legacy IPv4 domains - such as an H323 VoIP network.

Outstanding Issues in R5

There are still many issues that need to be settled before R5 standardisation is complete:

- Billing architecture.
- Interfaces to the HSS.
- Details of interfaces to application servers.
- Information flows for local and emergency services.
- Number and name portability issues.
- Mechanisms for interaction with visited network resources.
- Security solution and protocols to access the HSS.
- CS domain interworking.

9.4.3 Is R4/5 Worthy of the Term 'all IP'?

How does the R5 architecture compare with the all-IP design outlined earlier? On the face of it, it seems to fall a long way short:

- There is no native IP transport. The user data IP headers are never read or used for routing - these are still encapsulated within GTP tunnels from the RNC to the SGSN and from the SGSN to the GGSN. Web caching and multicast are not possible when tunnels are used for each route.
- There is no easy integration path for other access technologies, such as WLANs (which will be considered further later in this chapter).

- Security/HSS access from the SGSN and terminal is still via a separate SS7 signalling network.
- Session control uses SIP but does not follow the IP model of value chain separation, e.g. a user's network provider is still their service provider who is still their content provider through the Internet Multimedia domain, and when a user roams, their service access is controlled by their home network. In addition, 3GPP does not use SIP in the way that an IETF architecture would, e.g. security is done on the registration messages and not on the INVITEs, and registration must be complete before issuing an INVITE.
- The RAN is unchanged from the original version of UMTS except, perhaps, for supporting real-time packet handover. It is still an AAL2 switching cloud.
- A special bridge is needed to emulate multicast. Each terminal requires an IP tunnel to the bridge. A future anycast application would have to stop at the GGSN, for example.
- The network design still violates the end-to-end and layering IP principles. There is no support for end-to-end session creation - it takes place in the IM domain.
- The VHE functionality is still provided by the HSS and limited to services such as IP pre-pay. It does not allow access without the SIM card and does not support personal mobility.

9.4.4 CDMA2000 Evolution

Evolution is also planned of the original CDMA2000 1X standard, which was discussed in [Chapter 2](#). This is imaginatively called 'CDMA 1xEV' - EV stands for evolution, and 1X refers to the system using a single 1.25-MHz carrier. There is a two-phase strategy:

- CDMA 1xEV-DO - data only - which will increase the peak data rate on the down link to a (theoretical) 2.4 Mb/s, with perhaps 600 kb/s average. The DO service will have to be run on a separate carrier to basic CDMA2000 1X. This has now been recognized by the ITU as part of the 3G-IMT2000 family.
- CDMA 1xEV-DV - data and voice - which will once again allow voice and high speed data to operate on the same carrier. It will also enable real-time packet services and improved mechanisms for quality of service. It may also increase the bit rate further. 1x-DO is expected to become available for operators during 2002, whereas 1x-DV will be available about two years later.

The original CDMA evolution plans for wideband operation over three carriers (i.e. 3.75 MHz) - called 3X - has been abandoned for the moment, in favour of increasing the data rate within a single 1.25-MHz carrier.

Work is also going on to evolve CDMA2000's core network in the 3GPP2 standardisation group. It is developing an 'all IP' solution - sometimes called the NGN, next generation network. The path planned follows a similar rationale to UMTS evolution in R4 and R5 separating out the data and control planes by introducing a Media Gateway, Media Gateway Controller, and Signalling Gateway, and introducing the Megaco/H.248 and SIP protocols. A detailed view of the functional architecture is shown in [Figure 9.7](#).

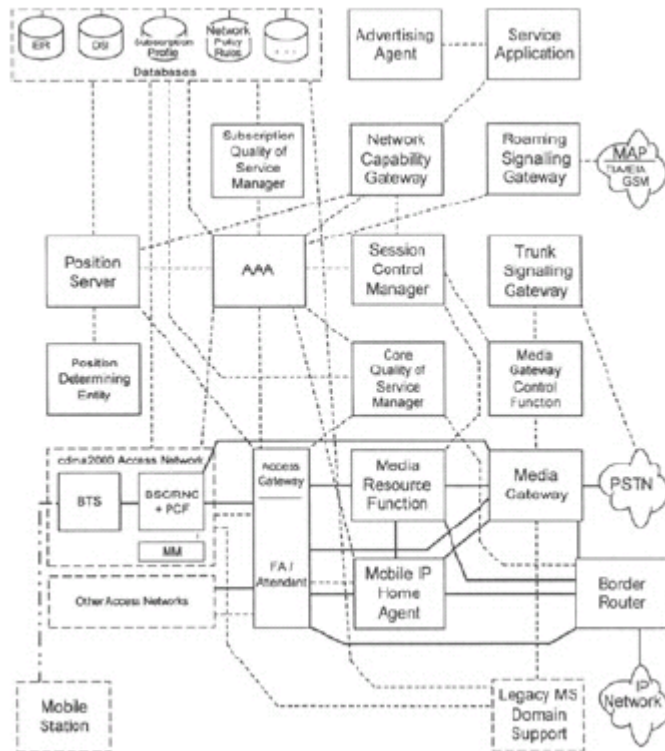


Figure 9.7: 3GPP2's 'all IP' reference architecture - a functional view.


CDMA2000 and UMTS core network evolutions might be expected to converge, although backward compatibility with their earlier incarnations is one stumbling block.

9.5 UMTS Beyond R5

There are a number of proposals for developments beyond R5 - not yet called R6 - from 3GIP and 3GPP:

- Mobile IP (MIP) - Proposed to replace the GTP tunnels from the GGSN to the RNC.
- EMA - Edge Mobility Architecture (MER-TORA as described in the mobility [Chapter7](#)) from the RNC to the GGSN. See [Figure 9.8](#).
- GTP to the Radio Network Controller (RNC) - A single GTP tunnel to the RNC from the GGSN -taking the SGSN out of the data path and turning it into an SGGN-server that only receives signalling messages. This would allow the SGSN control functions to be separately dimensioned from the user traffic.
- SGSN server- This is a proposal to split the SGSN into a server for the control parts and move the routing/network layer to a Media Gateway (MGW). This is similar to the R4 splitting of the MSC into a MGW and MGW controller.
- MGW/GSN Server - This is a proposal to have a MGW and Signalling Gateway (SGW) for each type of access technology in order to harmonise to a standard IP core.

Each pair of gateways adapts the data and signalling to fit with an IP core. Thus, the UTRAN would have a pair of gateways - as would a wireless LAN network.



Introducing MGWs would seem to be a regressive step from an IP point of view the whole point of IP is not to use gateways for signalling or traffic - except to legacy networks. The reason for the proposal to split the SGSN is primarily to allow separate signalling and data dimensioning and not really related to IP issues. Only the first two proposals would result in native IP transport down to the RNC. Mobile IP is already used in cdma2000 but has well-known difficulties: triangular routing and lack of real-time handover as well as security holes. These could be developed with fast IPv6 MIP handover, for example. The alternative seems to be MER-TORA, which would introduce native IP transport down to the RNC and allow a number of the advantages of IP (multicast to the RNC, caching, addition of WLANs, etc.). At the present time, however, 3GPP has not clarified what form post-R5 standards will take and seems likely to take a cautious approach until it has operational experience of R3.

Overall, these proposals are beginning to address the need for native IP routing in the PS domain and the need for the SGSN control functions to be separated from its routing functions. This could be R6 phase one - and, finally, for the routing functions to be replaced by Mobile IP or MERTORA as a second phase.

9.6 Wireless LANs

Many people think that wireless LANs have an important part to play in the future of 3G evolution because:

- They have a large amount of licence-exempt spectrum. Existing WLANs - such as 802.11b operate in the 2.4 GHz so-called ISM (Industrial Scientific and Medical) band. 80 MHz of spectrum in Europe (and a different allocation in the US) can be used without a specific licence. However, many countries (including the UK) do not, at present, allow public telecommunications to be offered in this spectrum, although this may change in the near future. More spectrum is potentially available to the nextgeneration WLANs - HIPERLAN/2 and 802.11 a - in the 5-GHz band (100 MHz+).
- They offer high bandwidths: 802.11b systems today can provide real user throughputs in excess of several Mbit/s.
- The cards are cheap - typically \$100 or so.
- Wireless LANs handle asymmetric traffic well.

The idea as already seen with Mary at the University is that WLANs could provide hotspot coverage (cafes, campuses, railway stations, offices, etc.). They could provide higher bandwidths than cellular systems at a lower cost (since the spectrum is free, and the cards are already high-volume low-cost items).

One of the major debates within the industry is how to 'unify' WLANs with 3G systems such as UMTS. One solution is very similar to our design for an all-IP mobile network – called loose coupling, the two systems only really share the same IP core network and a common authentication mechanism - either a user's SIM card authenticates that user, or they use a normal dial-IP password and identifier (NAI). Loose coupling is shown in [Figure 9.9](#). The alternative unification route being considered is tight coupling - where the WLANs provide only a link layer service, and all control/network functions are based on UMTS protocols (such as RANAP and GTP). ETSI BRAN is considering both loose and tight coupling for HIPERLAN-UMTS interworking.

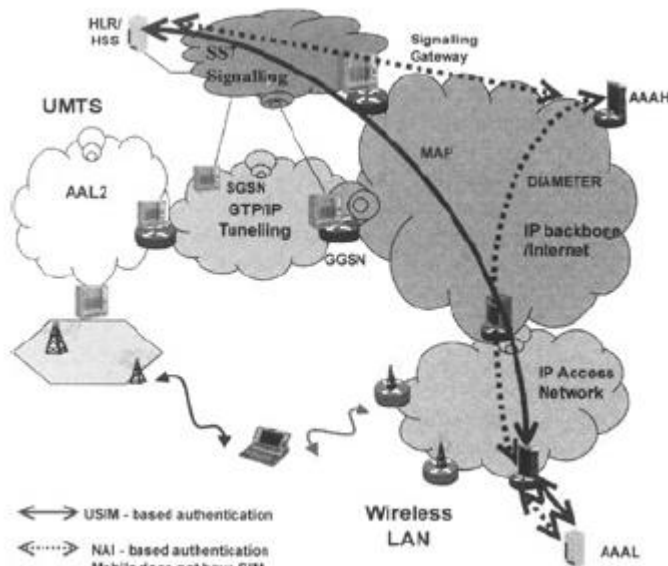


Figure 9.9: UMTS-HLAN Loose coupling proposal.

However, a question mark currently hangs over the next generation of WLANs, with worries about the cost of radio transceivers in the 5-GHz band and arguments over the maximum permitted power levels to avoid interference with satellite users of the 5-GHz band. The huge success and low cost of 802.11b is also pushing HIPERLAN products further into the future. It has just been announced that BT Group will launch a public WLAN service in the UK over the summer of 2002 - giving access to the Internet and corporate Intranets. Initially 20 sites are planned in airports, cafes and railway stations, rising to 400 hotspots after the first year and 4000 in three years.